UNI KASSEL
VERSITÄT

MASTER THESIS

Convergence of Modal Fixpoint Formulas over Special Classes of Structures

A thesis submitted in fulfillment of the requirements for the degree of Master of Science in the

Theoretical Computer Science / Formal Methods Group

at the

Department of Electrical Engineering and Computer Science

Author: Marco Sälzer

Filing Date: 07.11.2019

Submission Date: 16.04.2020

First Examiner: Prof. Dr. Martin Lange

Second Examiner: Prof. Dr. Stefan Göller

Supervisor: Dipl.-Math. Florian Bruse

Declaration of Authorship

I, Marco Sälzer, declare that this thesis titled, "Convergence of Modal Fixpoint Formulas over Special Classes of Structures" and the work presented in it are my own. I confirm that:

- This work was done wholly or mainly while in candidature for a research degree at the University of Kassel.
- Where any part of this thesis has previously been submitted for a degree or any other qualification at the University of Kassel or any other institution, this has been clearly stated.
- Where I have consulted the published work of others, this is always clearly attributed.
- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work.
- I have acknowledged all main sources of help.
- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself.

Place, Date, Signature

Acknowledgements

First of all, I would like to thank all persons dear to me. Especially those who have given me moral support over the past couple of months. I know it took a lot of patience, but i can assure you that this work would not exist without you.

Individually, I would like to thank the following persons: Prof. Martin Lange, for giving me the chance to work at this topic and for always giving the right insights at the right time. Prof. Stefan Göller, for agreeing to be the second examiner of this thesis. Lastly, i would like to give special thanks to my supervisor Florian Bruse. I cannot stress enough how important your guidance was throughout this time. I hope that a simple thank you can express how helpful our hours of discussion were for me.

Contents

List of Figures

Chapter 1

Introduction

Fixpoint logics play a major role in the area of formal logics in computer science and can be regarded as a classical extension of other logics like first-order, propositional or modal logic. The introduction of fixpoint operators to such logics enables the expression of recursive properties, which is of fundamental interest in many fields of formal computer science.

A well-known fixpoint logic is the Modal μ-Calculus ($\mathcal{L}_\mu$), usually considered in the form given by Kozen [Koz83]. It possesses high expressive power and incorporates other widely-studied logics like CTL* [Dam94]. Nevertheless, properties expressible by a $\mathcal{L}_\mu$-formula are at most regular [JW96]. This gives room for modal logics with the ability to express properties beyond regularity.

The Higher-Order Fixpoint Logic (HFL) is such a modal logic, which was first introduced by Viswanathan and Viswanathan [VV04]. It subsumes a simply-typed λ-calculus into the modal μ-calculus. From a semantic perspective this enables a formula of HFL to describe higher-order objects, e.g. functions, besides base-type objects like sets. The model-checking problem of HFL is still decidable, but due to the expressive power gained the problem is k-EXPTIME complete, where k corresponds to the highest order of used types [ALS07].

Ususally, fixpoint formulas of $\mathcal{L}_\mu$ (or HFL) are interpreted over Labeled Transition Systems (LTS). Then, Kleene‘s fixpoint theorem states that their semantics can be computed in an iterative fashion, which is not necessarily a finite iteration. Such an iterative procedure leads to various questions regarding convergence. Examples here are: ”What is the point of convergence for a specific fixpoint formula over a specific LTS?” or ”Is finite convergence guaranteed for a fragment of (fixpoint) formulas over some class of LTS?” and many more. These question belong to the recently studied field of convergence behaviour of fixpoint formulas, definable in the modal μ-calculus and modal logics in general. A short overview of results in this area is given in Sec. 1.1.

This work is concerned with the question whether the gap in expressive power of $\mathcal{L}_\mu$ and HFL comes with a gap in the convergence behaviour of fixpoint formulas of $\mathcal{L}_\mu$ and fixpoint formulas of HFL. A family of infinite words $w_{n,m}$ is investigated, which witness a difference in the convergence behaviour of $\mathcal{L}_\mu$ and HFL. The interpretation of $\mathcal{L}_\mu$-formulas and HFL-formulas over infinite words is done with linear-time variants $\mathcal{L}_\mu^{\text{lin}}$ and HFL_{lin}. Then, the difference between $\mathcal{L}_\mu$ and HFL is described as follows: It is suspected that the iteration of all fixpoints definable by a $\mathcal{L}_\mu^{\text{lin}}$-formula does converge after a finite number of steps over all $w_{n,m}$ and that there are fixpoints definable by an HFL_{lin}-formula such that its iteration does not converge after a finite number of steps for some $w_{n,m}$. To formally capture this, a suitable convergence criterion, called the finite convergence, is defined for a pair formula and structure. The iteration of a fixpoint formula is done in a syntactic way in this work, which is called unfolding a fixpoint formula. Then, a formula is said to be finitely converging over some structure if there is some n-th unfolding of all fixpoint subformulas, starting with the innermost and ending with the outermost, such that this unfolding equals the true semantics. Futhermore, it is required that all unfolded subformulas are equal to their respective true semantics as well. Such an n-th unfolding is semantically equal to a point of full stability of a fixpoint iteration.

In the main part of this work it is shown that all alternation-free $\mathcal{L}_\mu^{\text{lin}}$-formulas are finitely converging over all $w_{n,m}$. Additionally, a strong conjecture that this also implies the finite convergence of arbitrary $\mathcal{L}_\mu^{\text{lin}}$-calculus formulas over all $w_{n,m}$ is discussed. In contrast to this, this work presents rather simple examples of HFL_{lin}-formulas, which are not finitely converging for some $w_{n,m}$. The work is divided into the following parts:

- Chap. 2 defines all needed preliminaries. This includes basic syntactical and semantical concepts as well as $\mathcal{L}_\mu$ and the HFL with needed, existing results. Furthermore, the simultaneous bottom-up unfolding and the finite convergence of a formula are defined. A major part of the proofs in this work uses automata-theoretic arguments. To do this, a model of Alternating Parity Automata (APA) is defined and the well-known equivalence between $\mathcal{L}_\mu$ and APA is established. In the end, linear-time equivalents $\mathcal{L}_\mu^{\text{lin}}$, HFL_{lin} and a corresponding APA interpretation are introduced.
- Chap. 3 introduces the word-family $w_{n,m}$. Besides the plain introduction, needed properties of this family are discussed and formally captured.
- Chap. 4 proves that all alternation-free $\mathcal{L}_\mu^{\text{lin}}$-formulas are finitely converging over $w_{n,m}$. The proof is divided into two parts: First, the finite convergence is proven for all $\mathcal{L}_\mu^{\text{lin}}$-formulas with just one kind of fixpoint operator, which are called unipolar formulas. Then, this result is used to prove the finite convergence of all alternation-free formulas. At the end of this chapter, a strong conjecture that this also implies the general result is discussed. A formal proof lies beyond the scope of this thesis.
- Chap. 5 gives examples of HFL_{lin}-formulas without finite convergence for some instances of $w_{n,m}$. The presented formulas are from the first- and second-order fragment of HFL_{lin}, which correlates with the highest order of used types.
- Chap. 6 summarizes this works findings and discusses its implications for $\mathcal{L}_\mu$ and HFL. Furthermore, possible continuations of this work are presented.

1.1 Related Work

The main part of this work is concerned with the convergence of fixpoints definable in the modal μ-calculus. One major line of research in this area is the problem of closure ordinals, which is the least number of iterations a fixpoint formula needs to converge across all structures. More precisely, the question asked is: “For which ordinals does a formula exist with precisely this ordinal as closure ordinal?” There are several partial results. Czarnecki [Cza10] has given a construction principle for every ordinal below ω^2. Besides this result, an important contribution of Czarnecki's work to the problem of closure ordinals was the introduction of designated propositions, which in combination lift the closure ordinal of a formula. Afshari and Leigh [AL13] generalized this idea and proved that w^2 is a tight upper bound for closure ordinals of the alternation-free fragment of the modal μ-calculus. Another generalization of Czarnecki's idea was given by Gouveia and Santocanale [GS19]. In this work it was proven that ω_1 is a closure ordinal and that the set of closure ordinals is closed under ordinal sum. The problem of closure ordinals for the two-way modal μ-calculus was investigated by Milanese and Venema [MV19]. This work proves that all ordinals below ω^ω are closure ordinals in the two-way setting.

Considering the question of closure ordinals as a decision problem, there is the well-known result of Otto [Ott99] stating that it is decidable whether a modal μ-calculus formula is equivalent to a formula in (basic) modal logic. A direct implication is that the question if a formula has a finite closure ordinal is decidable as well. Furthermore, it was shown by Blumensath, Otto and Weyer [BOW14] that it is decidable for MSO-formulas if there is a finite closure ordinal over the class of infinite trees. Besides the fact that bisimulation-invariant MSO coincides with the modal μ-calculus [JW96] over transition systems, this work is of interest as it makes use of special weighted automata, lately called distance parity automata [CL08]. This model offers an interesting concept to argue about approximative semantics.

Chapter 2

Preliminaries

This chapter defines fundamentals, needed in the context of this work. Sec. 2.1 introduces basic structures like words, trees, transition systems and linear-time structures. The Modal μ-Calculus ($\mathcal{L}_\mu$) and related definitions are given in Sec. 2.2. This includes its syntax, its semantics, an alternation hierachy, the unfolding of fixpoint formulas, the concept of finite convergence and guardedness of $\mathcal{L}_\mu$-formulas. Sec. 2.3 presents the Higher-Order Fixpoint Logic (HFL) in sufficient detail. This includes its type system, syntax and semantics. Furthermore, a hierachy of fragments HFL^k is defined, depending on the highest order of used types, and a connection between $\mathcal{L}_\mu$ and the base-type fragment HFL^0 is established. In Sec. 2.4 is introduced a model of Alternating Parity Automata (APA) and the semantic equivalence of $\mathcal{L}_\mu$-formulas and APAs is recapped. The given APA model is tailored to this purpose. Sec. 2.5 presents linear-time variants for $\mathcal{L}_\mu$, HFL and APAs, which allow arguing about these concepts on infinite words instead of (linear-time) structures.

2.1 Structures

2.1.1 Words and Trees

Throughout this work, an alphabet is assumed to be finite and usually referred to by Σ.

Definition 2.1.1. Let Σ be an alphabet. A finite word w over Σ is a sequence $x_0\, x_1 \cdots x_{k-1}$ with $k \in \mathbb{N}$ such that $x_i \in \Sigma$ for all $i \in \{0, \ldots, k-1\}$. An infinite word w over Σ is an infinite sequence $x_0\, x_1 \cdots$ such that $x_i \in \Sigma$ for all $i \geq 0$.

The set of all finite words over some alphabet Σ is denoted by Σ^* and the set of all infinite words by Σ^ω. The set of prefixes $\text{pre}(w)$, the set of infixes $\text{inf}(w)$ and the set of suffixes $\text{suff}(w)$ of w are defined as

- $\text{pre}(w) = \{v \in \Sigma^* \mid \text{there is } u \in \Sigma^\omega \text{ such that } vu = w\}$,
- $\text{inf}(w) = \{v \in \Sigma^* \mid \text{there are } u \in \Sigma^*, u' \in \Sigma^\omega \text{ such that } uvu' = w\}$,
- $\text{suff}(w) = \{v \in \Sigma^\omega \mid \text{there is } u \in \Sigma^* \text{ such that } uv = w\}$,

where the concatenation of two words u and v is denoted by uv. Note, this concatenation is not defined if $u \in \Sigma^\omega$. Let $w = x_0\, x_1 \cdots$. The i-th letter x_i is addressed by $w(i)$, the prefix $x_0 \cdots x_i$ is denoted by $w[: i]$, the infix $x_i \cdots x_j$ by $w[i : j]$ and the suffix $x_i\, x_{i+1} \cdots$ by $w[i :]$. If $w \in \Sigma^*$ these definitions are given analogously.

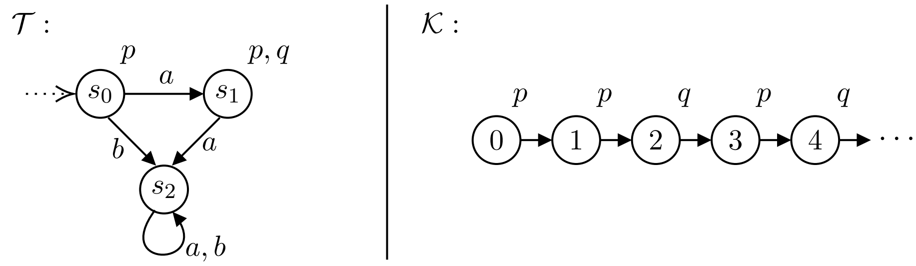

FIGURE 2.1: Visualisation of two LTS $(\mathcal{T}, s_0)$ and $\mathcal{K}$. The LTS $\mathcal{T}$ shows several relations and propositions and $\mathcal{K}$ is in fact a linTS.

Definition 2.1.2. Let Σ be an alphabet. A Σ-labeled tree is a partial function $t : \mathbb{N}^* \to \Sigma$, which satisfies the following conditions:

- t is prefix-closed, which means that if t is defined for some $v \in \mathbb{N}^*$ and $u \in \mathbb{N}^*$ is a prefix of v then t is defined for u.
- t is left-closed, which means that if t is defined for some $vi \in \mathbb{N}^*$ and $i > 0$ then t is defined for $v(i-1)$.

If $v \in \mathbb{N}^*$ is defined in a tree t it is called a node of t. For every non-empty tree it holds that $\epsilon \in \mathbb{N}^*$ is a node, which is called the root. Some node $vi \in \mathbb{N}^*$ is called a child of v and v is called the parent of vi. A branch in a tree t is a possibly infinite sequence $t(v_0)\, t(v_1) \cdots$ such that $v_0 = \epsilon$ and for all v_{i+1} with $i \geq 0$ it holds that $v_{i+1} = v_i j$ such that $v_i j$ is a node in t. Thus, a branch of t can be interpreted as an (infinite) word over the alphabet Σ. All nodes $u, v \in \mathbb{N}^*$ in a tree are said to be on one level if $|u| = |v|$. For the length $i \in \mathbb{N}$ the corresponding level is called the i-th level of ρ. A finite tree ρ is said to be of height k if the maximum level of ρ is the k-th. This is denoted by height(k).

2.1.2 Labeled Transition Systems

Definition 2.1.3. Let A, P be countable sets of symbols. A Labeled Transition System (LTS) is a tuple $\mathcal{T} = (S, \{\xrightarrow{a} \mid a \in A\}, L)$, where S is a finite set of states, each $\xrightarrow{a}$ describes a binary relation $\xrightarrow{a} \subseteq S \times S$ of transitions, and L is a labeling function of type $S \to 2^P$, which assigns a set of propositions to each state.

Let $\mathcal{T}$ be an LTS. The abbreviation $s \xrightarrow{a} t$ denotes $(s, t) \in \xrightarrow{a}$ and t is called an a-successor of s respectively s an a-predecessor of t. Sometimes $\mathcal{T}$ is considered with some designated state $s_0 \in S$, called the initial state. This is denoted by $(\mathcal{T}, s_0)$.

Usually, an LTS is visualized as given in Fig. 2.1. States are depicted as circles and relations as directed edges. If states are named, the name is written inside the circle. For the LTS $\mathcal{T}$ it holds that $S = \{s_0, s_1, s_2\}$. Valid propositions are written outside the circle. For instance, it holds that $L(s_1) = \{p, q\}$. Edges are named as well, to clarify to which relation they belong. The LTS $\mathcal{T}$ includes two relations $\xrightarrow{a}$ and $\xrightarrow{b}$. An initial state, if it exists, is denoted by an unique directed edge with no starting point. In the case of $\mathcal{T}$ it is the state s_0.

A special class of LTS are the so called linear-time structures.

Definition 2.1.4. Let P be a countable set of propositions. A Linear-Time Structure (linTS) is a special kind of labeled transition system defined as $(\mathbb{N}, \{(i, i+1) \mid i \in \mathbb{N}\}, \mathcal{K})$, where the state set is the set of natural numbers $\mathbb{N}$, the only relation is the usual successor relation on $\mathbb{N}$ and $\mathcal{K}$ is a labeling function $\mathbb{N} \to 2^P$.

A linear-time structure is completely defined by its labeling function $\mathcal{K}$. Therefore, a linTS is referred to by its labeling function $\mathcal{K}$ throughout this work. As for the case of general LTS, the tuple $(\mathcal{K}, i)$ with $i \in \mathbb{N}$ denotes a linear-time structure with specified initial state i. An example is given in Fig. 2.1 in the form of the linTS $\mathcal{K}$.

Remark 2.1.5. Let $w \in \Sigma^\omega$ be an infinite word. The corresponding linear-time structure $\mathcal{K}_w$ is defined as the function $\mathcal{K}_w : \mathbb{N} \to \Sigma$ with $i \mapsto x_i$, where x_i is the i-th letter in w and Σ is interpreted as a set of propositions.

Let $\mathcal{K} : \mathbb{N} \to 2^P$ be a linear-time structure. The corresponding infinite word $w_\mathcal{K}$ is defined as the infinite sequence $\mathcal{K}(0)\,\mathcal{K}(1)\,\cdots$ over the alphabet 2^P.

2.2 The Modal μ-Calculus

2.2.1 Syntax

Definition 2.2.1. Let P be a countable set of propositions, A be a countable set of actions and V be a countable set of variables. A formula φ is a modal μ-calculus formula ($\varphi \in \mathcal{L}_\mu$) if it is producible by the following grammar:

$$\varphi ::= p \mid \neg p \mid x \mid \bot \mid \top \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \langle a \rangle \varphi \mid [a]\varphi \mid \mu x.\, \varphi \mid \nu x.\, \varphi,$$

where $p \in P$, $a \in A$ and $x \in V$.

Throughout this work it is assumed that a variable x occurring in a formula φ is bound by μ or ν at most once and every bound variable in φ occurs in the scope of its binding operator only. This property is called well-namedness. A variable is called free in φ if it occurs unbound and φ is called a sentence if no free variables occur in φ. Let $\varphi, \psi \in \mathcal{L}_\mu$ be formulas. The abbreviation $\varphi[\psi/x]$ denotes the $\mathcal{L}_\mu$-formula which is obtained by simultaneously replacing every free occurrence of x in φ with ψ. The set of subformulas $\mathrm{sub}(\varphi)$ of φ is defined recursively as

$$\mathrm{sub}(\varphi) = \begin{cases} \{\varphi\} & \text{if } \varphi = p, x, \bot, \top, \\ \{\varphi\} \cup \{p\} & \text{if } \varphi = \neg p, \\ \{\varphi\} \cup \mathrm{sub}(\psi_1) \cup \mathrm{sub}(\psi_2) & \text{if } \varphi = \psi_1 \vee \psi_2, \psi_1 \wedge \psi_2, \\ \{\varphi\} \cup \mathrm{sub}(\psi) & \text{if } \varphi = \langle a \rangle \psi, [a]\psi, \\ \{\varphi\} \cup \mathrm{sub}(\psi) & \text{if } \varphi = \mu x.\, \psi, \nu x.\, \psi. \end{cases}$$

The size of a formula $|\varphi|$ is the amount of distinct subformulas, formally defined as $|\varphi| = |\mathrm{sub}(\varphi)|$.

The modal-depth of a formula φ is denoted by $\mathrm{md}(\varphi)$ and is defined as

$$\mathrm{md}(\varphi) = \begin{cases} 0 & \text{if } \varphi = p, x, \bot, \top, \\ 1 + \mathrm{md}(\psi) & \text{if } \varphi = \langle a \rangle \psi, [a]\psi, \\ \max(\mathrm{md}(\psi_1), \mathrm{md}(\psi_2)) & \text{if } \varphi = \psi_1 \vee \psi_2, \psi_1 \wedge \psi_2, \\ \mathrm{md}(\psi) & \text{if } \varphi = \sigma x.\, \psi. \end{cases}$$

The abbreviation $\sigma x.\, \psi$ denotes a $\mathcal{L}_\mu$-formula where $\sigma \in \{\nu, \mu\}$. Given a fixpoint variable x, which is bound in some formula φ, its fixpoint formula $\sigma x.\, \psi \in \mathrm{sub}(\varphi)$ is denoted by $\mathrm{fp}_\varphi(x)$. A formula φ is called a σ-unipolar if for all $\sigma' x.\, \psi \in \mathrm{sub}(\varphi)$ it holds that $\sigma' = \sigma$. If the actual σ is not relevant φ is simply called unipolar. The (fixpoint) nesting-depth of a formula φ is denoted by $\mathrm{nd}(\varphi)$ and defined recursively as

$$\mathrm{nd}(\varphi) = \begin{cases} 0 & \text{if } \varphi = p, x, \bot, \top, \langle a \rangle \psi, [a]\psi, \\ \max(\mathrm{nd}(\psi_1), \mathrm{nd}(\psi_2)) & \text{if } \varphi = \psi_1 \vee \psi_2, \psi_1 \wedge \psi_2, \\ 1 + \mathrm{nd}(\psi) & \text{if } \varphi = \sigma x.\, \psi. \end{cases}$$

2.2.2 Alternation Hierarchy

The set of $\mathcal{L}_\mu$-formulas defines an infinite hierachy, expressing the degree of nesting of least and greatest fixpoints. The following definition was first introduced by Niwiński [Niw86]. There are other alternation hierarchy definitions, for instance the hierarchy given by Emerson, Lei [EL86]. However, Niwiński's definition is used commonly becuase it better captures the interdependence between least and greatest fixpoints.

Definition 2.2.2. Let φ be a $\mathcal{L}_\mu$-formula. The alternation depth of φ is the least n ($\mathrm{ad}(\varphi) = n$) such that $\varphi \in \Sigma_{n+1} \cap \Pi_{n+1}$. For all $n \in \mathbb{N}$ the sets Σ_n and Π_n are inductively defined:

- Σ_0, Π_0 are the sets of $\mathcal{L}_\mu$-formulas such that for every $\varphi \in \Sigma_0$ and $\varphi \in \Pi_0$ it holds that: There is no formula $\sigma x.\, \psi$ for some variable x and $\psi \in \mathcal{L}_\mu$ in $\mathrm{sub}(\varphi)$. It holds that $\Sigma_0 = \Pi_0$ and besides formulas with free occurring variables these sets are equal to the set of modal logic formulas.
- Σ_{n+1} is defined as the least set that includes $\Sigma_n \cup \Pi_n$ and is closed under the following rules:
  1. If $\varphi, \psi \in \Sigma_{n+1}$ and $a \in A$ then $\varphi \vee \psi, \varphi \wedge \psi, \langle a \rangle \varphi, [a]\varphi \in \Sigma_{n+1}$.
  2. If $\varphi \in \Sigma_{n+1}$ and x is a variable then $\mu x.\, \varphi \in \Sigma_{n+1}$.
  3. If $\varphi, \psi \in \Sigma_{n+1}$ and $x \in V$ then $\varphi[\psi/x] \in \Sigma_{n+1}$ provided that no fixpoint operator in φ binds a free variable in ψ.
- Π_{n+1} is defined analogously. Instead of closure under the least fixpoint operator (2.) it demands closure under the greatest fixpoint operator.

A formula $\varphi \in \mathcal{L}_\mu$ is called alternation-free if $\mathrm{ad}(\varphi) \leq 1$. Informally, a formula φ is alternation-free if there is no fixpoint variable x which is bound by $\sigma x.\psi \in \mathrm{sub}(\varphi)$, such that x occurs in ψ' of a subformula $\sigma' x'.\psi' \in \mathrm{sub}(\varphi)$ such that $\sigma \neq \sigma'$. This excludes all formulas with mutually dependent least and greatest fixpoint formulas.

Example 2.2.3. Let $\varphi_1 = \mu x. (\nu y. p \wedge [b]y) \vee \langle a \rangle x$ and $\varphi_2 = \nu y. \mu x. (p \wedge \langle a \rangle y) \vee \langle a \rangle x$. Informally, φ_1 expresses the property "there is an a^*-successor such that from there on all b-paths satisfy p" and φ_2 expresses the property "there is an a-path on which there are infinitely many states where p holds". The formula φ_1 has an alternation-depth of 1 as there are two fixpoint-formulas in $\text{sub}(\varphi_1)$, but the variable of the outer does not occur in the scope of the inner. Thus, φ_1 is alternation-free. The formula φ_2, on the other hand, is a canonical example for a formula of an alternation-depth of 2. The reason for this is that the greatest fixpoint variable y occurs in the scope of the least fixpoint $\mu x. (p \wedge \langle a \rangle y) \vee \langle a \rangle x$.

2.2.3 Semantics

Definition 2.2.4. Let $\mathcal{T}$ be an LTS, η be a mapping $V \to 2^S$ called environment and φ a $\mathcal{L}_\mu$-formula. The semantics of φ regarding $\mathcal{T}$ and η, denoted by $[\![\varphi]\!]^{\mathcal{T}}_\eta$, are defined inductively:

$$
\begin{aligned}
&[\![p]\!]^{\mathcal{T}}_\eta := L(p) \qquad [\![\neg p]\!]^{\mathcal{T}}_\eta := S \setminus L(p) \\
&[\![x]\!]^{\mathcal{T}}_\eta := \eta(x) \qquad [\![\bot]\!]^{\mathcal{T}}_\eta := \emptyset \qquad [\![\top]\!]^{\mathcal{T}}_\eta := S \\
&[\![\psi_1 \vee \psi_2]\!]^{\mathcal{T}}_\eta := [\![\psi_1]\!]^{\mathcal{T}}_\eta \cup [\![\psi_2]\!]^{\mathcal{T}}_\eta \qquad [\![\psi_1 \wedge \psi_2]\!]^{\mathcal{T}}_\eta := [\![\psi_1]\!]^{\mathcal{T}}_\eta \cap [\![\psi_2]\!]^{\mathcal{T}}_\eta \\
&[\![\langle a \rangle \psi]\!]^{\mathcal{T}}_\eta := \{s \in S \mid \text{there is } t \in S \text{ such that } s \xrightarrow{a} t \text{ and } t \in [\![\psi]\!]^{\mathcal{T}}_\eta\} \\
&[\![[a]\psi]\!]^{\mathcal{T}}_\eta := \{s \in S \mid \text{for all } t \in S \text{ with } s \xrightarrow{a} t \text{ holds } t \in [\![\psi]\!]^{\mathcal{T}}_\eta\} \\
&[\![\mu x. \psi]\!]^{\mathcal{T}}_\eta := \bigcap \{U \subseteq 2^S \mid [\![\psi]\!]^{\mathcal{T}}_{\eta[x \mapsto U]} \subseteq U\} \\
&[\![\nu x. \psi]\!]^{\mathcal{T}}_\eta := \bigcup \{U \subseteq 2^S \mid [\![\psi]\!]^{\mathcal{T}}_{\eta[x \mapsto U]} \supseteq U\}.
\end{aligned}
$$

The update $\eta[x \mapsto U]$ of η maps x to U and everything else according to η.

Note that the powerset of a set forms a complete lattice. Following this, the Knaster-Tarski Theorem [Tar+55] ensures the existence of least and greatest fixpoints for monotonic formulas. As Def. 2.2.1 only allows negations in front of atomic propositions, the existence of fixpoints in the definition of $[\![\mu x. \psi]\!]^{\mathcal{T}}_\eta$ and $[\![\nu x. \psi]\!]^{\mathcal{T}}_\eta$ is guaranteed. Let $\mathcal{T}$ be an LTS. If a state s is an element of $[\![\varphi]\!]^{\mathcal{T}}_\eta$ it is said that s satisfies φ. A set of states $U \subseteq S$ is said to be definable by $\varphi \in \mathcal{L}_\mu$ if it holds that $[\![\varphi]\!]^{\mathcal{T}}_\eta = U$.

2.2.4 Fixpoint Unfolding and Finite Convergence

From Kleene's fixpoint theorem it is known that the semantics of a fixpoint definable by a $\mathcal{L}_\mu$-formula can be approximated iteratively. Ususally, these iterations are considered beyond finitely many steps. To do this, the concept of ordinal numbers is used.

Definition 2.2.5. Let $\sigma x.\psi \in \mathcal{L}_\mu$ be a fixpoint formula, $\mathcal{T}$ an LTS, η an environment, α a limit ordinal. The α-th approximation of the formula $[\![\sigma x.\psi]\!]^{\mathcal{T}}_\eta$ is denoted by

$[\![\sigma x.\psi,\alpha]\!]^{\mathcal{T}}_{\eta}$ and defined as

$$[\![\mu x.\psi,\alpha]\!]^{\mathcal{T}}_{\eta} := \begin{cases} \emptyset & \text{if } \alpha = 0, \\ [\![\psi]\!]^{\mathcal{T}}_{\eta[x\mapsto[\![\mu x.\psi,\beta]\!]^{\mathcal{T}}_{\eta}]} & \text{if } \alpha = \beta + 1, \\ \bigcup_{\beta<\alpha}[\![\sigma x.\psi,\beta]\!]^{\mathcal{T}}_{\eta} & \text{if } \alpha \text{ is a limit ordinal and} \end{cases}$$

$$[\![\nu x.\psi,\alpha]\!]^{\mathcal{T}}_{\eta} := \begin{cases} S & \text{if } \alpha = 0, \\ [\![\psi]\!]^{\mathcal{T}}_{\eta[x\mapsto[\![\mu x.\psi,\beta]\!]^{\mathcal{T}}_{\eta}]} & \text{if } \alpha = \beta + 1, \\ \bigcap_{\beta<\alpha}[\![\sigma x.\psi,\beta]\!]^{\mathcal{T}}_{\eta} & \text{if } \alpha \text{ is a limit ordinal.} \end{cases}$$

As noted in Sec. 2.2.3, every ψ of a fixpoint formula $\mu x.\,\psi$ can be seen as a monotonic function $\psi(x) : 2^P \to 2^P$, which leads to the convergence of the value of $[\![\sigma x.\,\psi,\alpha]\!]^{\mathcal{T}}_{\eta}$ with increasing α and thus for every fixpoint formula $\sigma x.\,\psi$, LTS $\mathcal{T}$ and η there is a least ordinal α such that $[\![\sigma x.\,\psi,\alpha]\!]^{\mathcal{T}}_{\eta} = [\![\sigma x.\,\psi,\alpha+1]\!]^{\mathcal{T}}_{\eta} = [\![\sigma x.\,\psi]\!]^{\mathcal{T}}_{\eta}$. If there is such least α across all LTS it is called the closure ordinal of $\sigma x.\,\psi$.

This notion of fixpoint iteration can be generalized in the sense that each fixpoint defined by subformula is iteratively approximated as well.[1] There is a semantic and a syntactic way to define this. As this work is only interested in iterations below ω, it is more convenient to use the syntactic variant and restricting the iteration to finite steps, which is then called the unfolding of a formula. Therefore, let $s = (c_0, c_1, \dots, c_{k-1})$ be a tuple such that $c_i \in \mathbb{N}$. An update of s is denoted by $s[d/c_i]$, which is defined as $(c_0, c_1, \dots, c_{i-1}, d, c_{i+1}, \dots, c_{k-1})$. The update $(s+i)$ is defined as the tuple $s + i = (c_0 + i, c_1 + i, \dots, c_{k-1} + i)$.

Definition 2.2.6. Let φ be a $\mathcal{L}_\mu$-formula and $\{x_0, \dots, x_{k-1}\}$ the set of all pairwise different fixpoint variables in φ. The formula φ^s denotes the an unfolding of φ specified by s, which is defined as

$$\begin{array}{lll} p^s := p, & (\neg p)^s := \neg p, & x^s := x, \\ \bot^s := \bot, & \top^s := \top, & (\psi_1 \vee \psi_2)^s := \psi_1^s \vee \psi_2^s, \\ (\psi_1 \wedge \psi_2)^s := \psi_1^s \wedge \psi_2^s, & (\langle a\rangle\psi)^s := \langle a\rangle(\psi^s), & ([a]\psi)^s := [a](\psi^s), \end{array}$$

$$(\sigma x_i.\,\psi_i)^s := \begin{cases} \bot & \text{if } c_i = 0 \text{ and } \sigma = \mu, \\ \top & \text{if } c_i = 0 \text{ and } \sigma = \nu, \\ (\psi_i[(\sigma x_i.\,\psi)^{s[c_i-1/c_i]}/x_i])^{s[0/c_i]} & \text{if } c_i > 0. \end{cases}$$

This formalizes the notion of simultaneously unfolding each fixpoint in φ from the innermost to the outermost. The bottom-up approach avoids the duplication of inner fixpoint formulas while unfolding outer ones. If it holds that $c_0 = c_1 = \dots = c_{k-1} = n$ the corresponding tuple $(c_0, \dots, c_{k-1})$ is simply denoted by n.

Example 2.2.7. Consider the formula $\varphi = \mu x_0.\,(\nu x_1.\,p \wedge [b]x_1) \vee \langle a\rangle x_0$. Apart from different designations, this is the same formula as φ_1 of Ex. 2.2.3. Per definition the unfolding φ^2 is given by

$$\Big(\big((\nu x_1.\,p \wedge [b]x_1) \vee \langle a\rangle x_0\big)\big[(\mu x_0.\,(\nu x_1.\,p \wedge [b]x_1) \vee \langle a\rangle x_0)^{(1,2)}/x_0\big]\Big)^{(0,2)}.$$

[1] It should be mentioned that the definition in [AL13] seems suitable for this purpose, but is flawed when considering nested fixpoints.

The two times unfolding of $(\nu x_1.\, p \wedge [b]x_1) \vee \langle a \rangle x_0$ is defined as $p \wedge [b](p \wedge [b]\top) \vee \langle a \rangle x_0$. This means that $(\mu x_0.\,(\nu x_1.\, p \wedge [b]x_1) \vee \langle a \rangle x_0)^{(1,2)} = p \wedge [b](p \wedge [b]\top) \vee \langle a \rangle \bot$. In conclusion φ^2 is equal to

$$p \wedge [b](p \wedge [b]\top) \vee \langle a \rangle (p \wedge [b](p \wedge [b]\top) \vee \langle a \rangle \bot).$$

It can be seen that a state satisfies this formula if it or a direct a-successor satisfies $p \wedge [b]p$. This corresponds to the notion of proceeding exactly two iterations of the least and two of the greatest fixpoint.

If a unipolar formula φ is considered then the monotonicity properties of a single fixpoint iteration are transferred to the unfolding.

Lemma 2.2.8. Let $\varphi \in \mathcal{L}_\mu$ be a σ-unipolar formula, $\mathcal{T}$ an LTS and η an environment. It holds that $[\![\varphi^0]\!]^{\mathcal{T}}_\eta \subseteq [\![\varphi^1]\!]^{\mathcal{T}}_\eta \subseteq \cdots \subseteq [\![\varphi]\!]^{\mathcal{T}}_\eta$ if $\sigma = \mu$ and $[\![\varphi^0]\!]^{\mathcal{T}}_\eta \supseteq [\![\varphi^1]\!]^{\mathcal{T}}_\eta \supseteq \cdots \supseteq [\![\varphi]\!]^{\mathcal{T}}_\eta$ if $\sigma = \nu$.

Proof. The proof is done for the $\sigma = \mu$ case. For the case $\sigma = \nu$ the proof works in the same way, using dual arguments. W.l.o.g. it is assumed that $\varphi = \mu x_0.\, \psi_0$ and that there are k distinct fixpoint subformulas in φ. The proof is done by induction on the nesting-depth of $\mu x_0.\, \psi_0$.

If $\mathrm{nd}(\mu x_0.\, \psi_0) = 1$ the statement is a result of the monotonicity of the usual iteration of a single, least fixpoint.

Assume that $\mathrm{nd}(\mu x_0.\, \psi_0) = n + 1$ and that the statement is valid for all formulas of nesting-depth up to n. Fix $c_1 = c_2 = \cdots = c_{k-1} = i$. By induction it is shown for each case of c_0 that $[\![(\mu x_0.\, \psi_0)^s]\!]^{\mathcal{T}}_\eta \subseteq [\![(\mu x_0.\, \psi_0)^{(s+1)}]\!]^{\mathcal{T}}_\eta$ and $[\![(\mu x_0.\, \psi_0)^s]\!]^{\mathcal{T}}_\eta \subseteq [\![\mu x_0.\, \psi_0]\!]^{\mathcal{T}}_\eta$. With $c_0 = i$ this implies the statement of the lemma for this case.

The case $c_0 = 0$ is straightforward. Assume that $s = (l + 1, c_1, \ldots, c_{k-1})$ and that $[\![(\mu x_0.\, \psi_0)^{s'}]\!]^{\mathcal{T}}_\eta \subseteq [\![(\mu x_0.\, \psi_0)^{(s'+1)}]\!]^{\mathcal{T}}_\eta$ and $[\![(\mu x_0.\, \psi_0)^{s'}]\!]^{\mathcal{T}}_\eta \subseteq [\![\mu x_0.\, \psi_0]\!]^{\mathcal{T}}_\eta$ hold for $s' = (l, c_1, \ldots, c_{k-1})$. It holds that $(\mu x_0.\, \psi_0)^s = (\psi_0[(\mu x_0.\, \psi_0)^{s[l/c_0]}/x_0])^{s[0/c_0]}$ and that $\mathrm{nd}(\psi_0[(\mu x_0.\, \psi_0)^{s[l/c_0]}/x_0]) \leq n$. Then, by induction on the nesting-depth it is implied that

$$[\![(\psi_0[(\mu x_0.\, \psi_0)^{s[l/c_0]}/x_0])^{s[0/c_0]}]\!]^{\mathcal{T}}_\eta \subseteq [\![(\psi_0[(\mu x_0.\, \psi_0)^{s[l/c_0]}/x_0])^{(s+1)[0/c_0]}]\!]^{\mathcal{T}}_\eta.$$

Note that ψ_0 is monotonic in x_0 and $s[l/c_0] = s'$. Thus, by induction on c_0 follows that

$$\begin{aligned}[\![(\psi_0[(\mu x_0.\, \psi_0)^{s[l/c_0]}/x_0])^{(s+1)[0/c_0]}]\!]^{\mathcal{T}}_\eta &\subseteq [\![(\psi_0[(\mu x_0.\, \psi_0)^{s'+1}/x_0])^{(s+1)[0/c_0]}]\!]^{\mathcal{T}}_\eta \\ &= [\![(\psi_0[(\mu x_0.\, \psi_0)^{(s+1)[l+1/c_0]}/x_0])^{(s+1)[0/c_0]}]\!]^{\mathcal{T}}_\eta,\end{aligned}$$

which is the same as $[\![(\mu x.\, \psi)^{s+1}]\!]^{\mathcal{T}}_\eta$. The statement $[\![(\mu x_0.\, \psi_0)^s]\!]^{\mathcal{T}}_\eta \subseteq [\![\mu x_0.\, \psi_0]\!]^{\mathcal{T}}_\eta$ is proven in the same way. □

Similar to the iteration of a single fixpoint, a stability criterion can be defined for the unfolding of a formula φ. As the unfolding also considers subformulas, it is not sufficient to solely demand semantic equivalence of φ and some unfolding φ^n. Additionally, the unfolding of a subformula must be equivalent to its true semantics as well, provided that free occurring fixpoint variables are mapped to their true semantics. For a subformula $\sigma x.\, \psi \in \mathrm{sub}(\varphi)$ of some $\varphi \in \mathcal{L}_\mu$ let η_x be the update of an

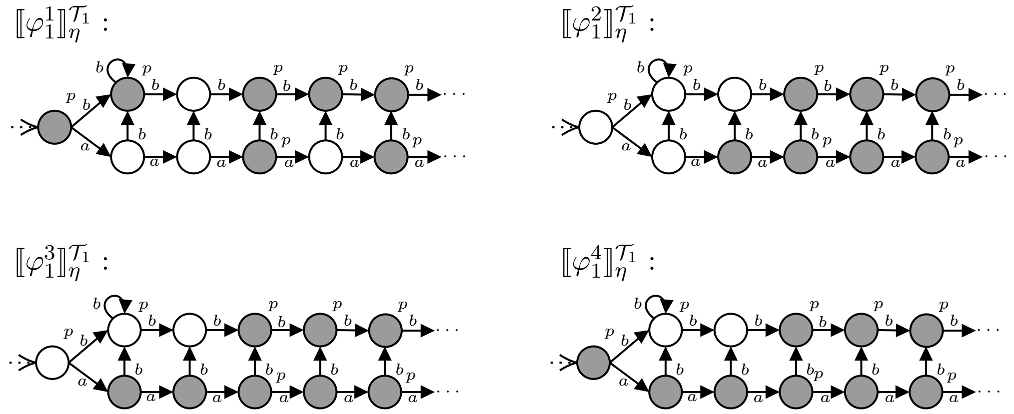

FIGURE 2.2: Visualisation of the semantics of unfoldings $\varphi_1^1, \varphi_1^2, \varphi_1^3$ and φ_1^4 of φ_1 of Ex. 2.2.10 over the given LTS $\mathcal{T}_1$. Filled states are included in the respective semantics, unfilled are not.

environment η, which maps every fixpoint variable $x', x' \neq x$ that occurs freely in ψ and to $[\![\sigma' x'.\psi']\!]_{\eta_{x'}}^{\mathcal{T}}$ with $\mathrm{fp}_\varphi(x') = \sigma' x'.\psi'$. If there is no such variable it holds that $\eta_x = \eta$.

Definition 2.2.9. Let $\varphi \in \mathcal{L}_\mu$ be a formula, $\mathcal{T}$ an LTS and η an environment. The formula φ is called finitely converging over $\mathcal{T}$ if there is n such that $[\![\varphi^n]\!]_\eta^{\mathcal{T}} = [\![\varphi]\!]_\eta^{\mathcal{T}}$ and for all $\sigma x.\, \psi \in \mathrm{sub}(\varphi)$ it holds that $[\![(\sigma x.\, \psi)^n]\!]_{\eta_x}^{\mathcal{T}} = [\![\sigma x.\, \psi]\!]_{\eta_x}^{\mathcal{T}}$. The set of all finitely converging $\mathcal{L}_\mu$-formulas over $\mathcal{T}$ is denoted by $\mathrm{FC}(\mathcal{T})$.

Informally, if $\varphi \in \mathrm{FC}(\mathcal{T})$ it means that there is n such that the maximum amount of unfoldings needed for any fixpoint subformula in φ to reach its true semantics is n times, provided that outer fixpoints are stable already.

Example 2.2.10. Consider $\varphi_1 = \mu x_0.\, (\nu x_1.\, p \wedge [b]x_1) \vee \langle a \rangle x_0$ of Ex. 2.2.7 again. The semantics of the first four unfoldings of φ_1 regarding LTS $\mathcal{T}_1$ are depicted in Fig. 2.2. The unfolding φ_1^0 is left out as its semantics are trivial. There are a two things to take from this example. Firstly, in general there are no monotonicity properties regarding the semantics of unfoldings. This can be seen in the difference of $[\![\varphi_1^1]\!]_\eta^{\mathcal{T}_1}$ and $[\![\varphi_1^2]\!]_\eta^{\mathcal{T}_1}$. The reason for this is the nesting of greatest and least fixpoint formulas in φ_1. Secondly, it can be seen that with unfolding φ_1^4 the semantics become stable. This is not a sufficient criteria for the finite convergence of φ_1, but with its simple form comes that φ_1^4 is also a point of finite convergence of φ_1 over $\mathcal{T}_1$.

Consider now the formula $\varphi_2 = \mu x.\, p \vee \langle a \rangle x$ and the LTS $\mathcal{T}_2$ in Fig. 2.3. The increasing distances between states with proposition p leads to $[\![\varphi_2^n]\!]_\eta^{\mathcal{T}_2} \neq [\![\varphi_2^{n+1}]\!]_\eta^{\mathcal{T}_2}$ for all $n \in \mathbb{N}$ This implies that φ_2 is not finitely converging over $\mathcal{T}_2$.

2.2.5 Guardedness

Definition 2.2.11. Let φ be a $\mathcal{L}_\mu$-formula. φ is called a guarded formula if every occurrence of a fixpoint variable $x \in \mathrm{sub}(\varphi)$ is in the scope of some modality operator $\langle a \rangle$ or $[a]$ which occurs inside $\mathrm{fp}_\psi(x)$.

Correspondingly, a fixpoint-variable x is called guarded if it occurs under the scope of some modality operator which occurs inside $\mathrm{fp}_\varphi(x)$.

It is well-known that for every $\mathcal{L}_\mu$-formula there is an equivalent guarded one [Wal00], but all known translations come with an exponential blowup in the size of the formula as shown in [BFL15]. The translation principles given by Kupferman et al. [KVW00] and Mateescu [Mat02] rely on two principles. These are presented here again as later given proofs make use of the transformation in detail.

The first principle is the observation that unfolding a fixpoint formula does not change its true semantics.

Proposition 2.2.12. Let $\sigma x.\, \psi \in \mathcal{L}_\mu$ be a fixpoint formula. For every LTS $\mathcal{T}$ and environment η it holds that $[\![\sigma x.\, \psi]\!]_\eta^{\mathcal{T}} = [\![\psi[(\sigma x.\, \psi)/x]]\!]_\eta^{\mathcal{T}}$.

The second principle needs some specifications: For a fixpoint formula $\sigma x.\, \psi \in \mathcal{L}_\mu$ let $\hat{\sigma}$ be $\bot$ if $\sigma = \mu$ and $\top$ if $\sigma = \nu$. An occurrence of a variable x is called weakly-guarded if it is is guarded or under the scope of another fixpoint quantifier which is inside $\mathrm{fp}_\varphi(x)$. The second principle is then given by the observation that replacing not weakly-guarded occurrences of x in a fixpoint formula $\sigma x.\, \psi$ with $\hat{\sigma}$ does not change its semantics.

Proposition 2.2.13. Let $\sigma x.\, \psi \in \mathcal{L}_\mu$ and let $\sigma x.\, \psi'$ be the result from replacing every occurrence of x which is not weakly-guarded in ψ by $\hat{\sigma}$. Then it holds for all LTS $\mathcal{T}$ and environments η that $[\![\sigma x.\, \psi]\!]_\eta^{\mathcal{T}} = [\![\sigma x.\, \psi']\!]_\eta^{\mathcal{T}}$.

An $\mathcal{L}_\mu$-formula can then be transformed into an equivalent guarded formula by applying these principles in a bottom-up approach to all fixpoint subformulas $\sigma x.\, \psi$. In each step lower fixpoint formulas are unrolled once and then remaining weakly-guarded occurrences of x are replaced with $\hat{\sigma}$. This ensures guardedness and does not change the semantics as seen in the propositions above.

Example 2.2.14. Consider the unguarded formula $\varphi = \nu x.\, (\mu y.\, (p \wedge x) \vee \langle a \rangle y) \wedge (\mu z.\, (q \wedge [a]x) \vee z) \wedge x$. The sole purpose of this formula is to demonstrate how the guarding procedure works. The first step in the guarding procedure considers innermost fixpoint subformulas and replaces all unguarded occurrences of their respective fixpoint variable with $\hat{\sigma}$. In the case of φ innermost fixpoint subformulas are $\mathrm{fp}_\varphi(y)$ and $\mathrm{fp}_\varphi(z)$. After the first step the formula φ_{tmp} is given by

$$\varphi_{\mathrm{tmp}} = \nu x.\, (\mu y.\, (p \wedge x) \vee \langle a \rangle y) \wedge (\mu z.\, (q \wedge [a]x) \vee \bot) \wedge x.$$

Nothing changed from $\mathrm{fp}_\varphi(y)$ to $\mathrm{fp}_{\varphi_{\mathrm{tmp}}}(y)$ as there is no unguarded occurrence of y in $\mathrm{fp}_\varphi(y)$. The unguarded occurrence of z in $\mathrm{fp}_\varphi(z)$ was replaced by $\bot$ in $\mathrm{fp}_{\varphi_{\mathrm{tmp}}}(z)$. Consider now the outer fixpoint formula $\mathrm{fp}_{\varphi_{\mathrm{tmp}}}(x)$. First, each inner fixpoint formula is unrolled once, which gives

$$\varphi_{\mathrm{tmp}_2} = \nu x.\, ((p \wedge x) \vee \langle a \rangle (\mu y.\, (p \wedge x) \vee \langle a \rangle y)) \wedge ((q \wedge [a]x) \vee \bot) \wedge x.$$

Two things are to be noted here: Firstly, the already ensured guardedness of all y occurrences in φ_{tmp} ensures that all $x \in \mathrm{fp}_{\varphi_{\mathrm{tmp}_2}}(y)$ are guarded. Secondly, as all occurrences of z vanished in the previous steps it follows that the subformula $\mathrm{fp}_{\varphi_{\mathrm{tmp}}}(z)$ is not present in φ_{tmp_2}. The last step in this example is to replace all not weakly-guarded occurrences of x in φ_{tmp_2}. This gives the guarded variant φ' of φ, which is defined as

$$\varphi' = \nu x.\, ((p \wedge \top) \vee \langle a \rangle (\mu y.\, (p \wedge x) \vee \langle a \rangle y)) \wedge ((q \wedge [a]x) \vee \bot) \wedge \top.$$

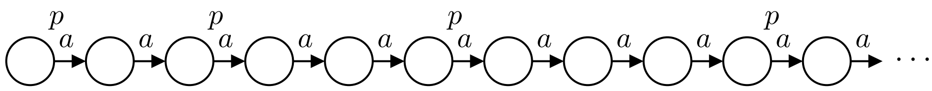

FIGURE 2.3: LTS $\mathcal{T}_2$ with increasing distances between two occurrences of p. The formula φ_2 of Ex. 2.2.10 is not finitely converging over this structure.

2.3 Higher-Order Fixpoint Logic

Higher-Order Fixpoint Logic (HFL) is an extension of the modal μ-calculus, incorporating a simply-typed λ-calculus. This lifts its expressional power to higher-order objects, but brings with it the need for a type system.

2.3.1 Higher-Order Types

Definition 2.3.1. τ is an HFL-type if it is producible by the following grammar:

$$\tau ::= \bullet \mid \tau^v \to \tau$$

where $v \in \{+, -, 0\}$ is called variance, $\bullet$ is the base-type. All other derivable types are called function-types.

If not further stated, a HFL-type is simply called type. For two given types τ, τ' the variance v in $\tau^v \to \tau'$ denotes whether the argument is used monotonically $(+)$, anti-monotonically $(-)$ or in an arbitrary way (0). The arrow operator $\to$ is right-associative, which means that each function type is isomorphic to one of the form $\tau_1 \to \cdots \to \tau_n \to \bullet$ for some $n \in \mathbb{N}$.

The order of a type is defined inductively: $\mathrm{ord}(\bullet) = 0, \mathrm{ord}(\tau^v \to \tau') = \max(\mathrm{ord}(\tau) + 1, \mathrm{ord}(\tau'))$.

Definition 2.3.2. Let $\mathcal{T}$ be a labeled transition system. The semantics $\mathcal{T}[\![\tau]\!]$ of a type τ over $\mathcal{T}$ is defined inductively as

$$\mathcal{T}[\![\bullet]\!] := (2^{\mathbb{N}}, \sqsubseteq_\bullet),$$
$$\mathcal{T}[\![\tau^v \to \tau']\!] := (\mathcal{T}[\![\tau]\!]^v \to \mathcal{T}[\![\tau']\!], \sqsubseteq_{\tau^v \to \tau'}),$$

where $(S, \leq)^+ = (S, \leq), (S, \leq)^- = (S, \geq)$ and $(S, \leq)^0 = (S, \leq)^+ \cap (S, \leq)^-$ for some order $\leq$. $\mathcal{T}[\![\tau]\!]^v \to \mathcal{T}[\![\tau']\!]$ is the set of monotonic functions from $\mathcal{T}[\![\tau]\!]^v$ to $\mathcal{T}[\![\tau']\!]$. The order $\sqsubseteq_\bullet$ denotes the usual set inclusion $\subseteq$, and $\sqsubseteq_{\tau^v \to \tau'}$ is the pointwise ordering of functions from $\mathcal{T}[\![\tau]\!]^v$ to $\mathcal{T}[\![\tau']\!]$: $f \sqsubseteq_{\tau^v \to \tau'} g$ iff for all $x \in (\mathcal{T}[\![\tau]\!])^v$ we have $f(x) \sqsubseteq_{\tau'} g(x)$.

It is a well-known fact that the power set of a set ordered by inclusion forms a complete lattice. Let $L = (S, \leq)$ be a complete lattice. Then L^+ and L^- of Def. 2.3.2 are complete lattices as well. Furthermore, in the way the semantics of a type are defined it holds that each function-type defines a complete lattice, too.

$$\frac{}{\Gamma \vdash p{:}\bullet} \qquad \frac{v \in \{+,0\}}{\Gamma, x^{v,\tau} \vdash x{:}\tau} \qquad \frac{\Gamma \vdash \varphi{:}\bullet}{\Gamma \vdash \langle a \rangle \varphi{:}\bullet} \qquad \frac{\Gamma^- \vdash \varphi{:}\bullet}{\Gamma \vdash \neg\varphi{:}\bullet} \qquad \frac{\Gamma \vdash \varphi_1{:}\bullet \qquad \Gamma \vdash \varphi_2{:}\bullet}{\Gamma \vdash \varphi_1 \vee \varphi_2{:}\bullet}$$

$$\frac{\Gamma, x^{v,\tau} \vdash \varphi{:}\tau'}{\Gamma \vdash \lambda(x^{v,\tau}).\, \varphi{:}\tau^v \to \tau'} \qquad \frac{\Gamma, x^{+,\tau} \vdash \varphi{:}\tau}{\Gamma \vdash \mu(x^\tau).\, \varphi{:}\tau} \qquad \frac{\Gamma \vdash \varphi{:}\tau^+ \to \tau' \qquad \Gamma \vdash \varphi'{:}\tau}{\Gamma \vdash \varphi\,\varphi'{:}\tau'}$$

$$\frac{\Gamma \vdash \varphi{:}\tau^- \to \tau' \qquad \Gamma^- \vdash \varphi'{:}\tau}{\Gamma \vdash \varphi\,\varphi'{:}\tau'} \qquad \frac{\Gamma \vdash \varphi{:}\tau^0 \to \tau' \qquad \Gamma \vdash \varphi'{:}\tau \qquad \Gamma^- \vdash \varphi'{:}\tau}{\Gamma \vdash \varphi\,\varphi'{:}\tau'}$$

FIGURE 2.4: The HFL typing rules.

2.3.2 Syntax

Definition 2.3.3. Let A be a set of transition names, P a set of propositions and V a set of variable names. φ is called a HFL-formula ($\varphi \in$ HFL) if it is producible by the following grammar:

$$\varphi ::= p \mid x \mid \neg\varphi \mid \varphi \vee \varphi \mid \langle a \rangle \varphi \mid \lambda(x^{v,\tau}).\, \varphi \mid \varphi\,\varphi \mid \mu(x^\tau).\, \varphi$$

where $p \in P$, $a \in A$, $x \in V$, τ is an HFL-type and v a variance.

In addition to this basic grammar, there are several usual abbreviations:

$$\begin{aligned} &\bot = p \wedge \neg p, && \top = p \vee \neg p, \\ &\varphi \wedge \psi = \neg(\neg\varphi \vee \neg\psi), && [a]\varphi = \neg\langle a \rangle(\neg\varphi), \\ &\nu x.\varphi = \neg\mu x.\neg\varphi[\neg x/x], && \end{aligned}$$

for some proposition $p \in P$, action $a \in A$ and variable $x \in X$. The set of subformulas $\mathrm{sub}(\varphi)$ of a HFL-formula is defined in the same manner as it is defined for $\mathcal{L}_\mu$-formulas in Sec. 2.2.1.

Not every derivable HFL-formula using this grammar is meaningful. For example an application $\psi_1\,\psi_2$ needs to respect the types of ψ_1 and ψ_2 in order to have meaningful semantics. Additionally, it must be ensured that every occurrence of the variable x in a formula $\mu(x^{v,\tau}).\varphi$ (as well as in the ν case) appears under an even number of negations to guarantee monotonicity and the existence of a fixpoint.

To ensure these properties, HFL makes use of its type system defined in Def. 2.3.1 and Def. 2.3.2. A sequence of $\Gamma = x_0^{v_0,\tau_0}, x_1^{v_1,\tau_1}, \ldots, x_{n-1}^{v_{n-1},\tau_{n-1}}$, where x_i is a variable, v_i a variance and τ_i a type, is called a context. It is assumed that these contexts are well-named, which means that no variable appears twice in a context. Γ^- is defined as the context where every $v_i = +$ is replaced by $v_i' = -$ and vice-versa.

An HFL-formula is called well-typed for some context Γ and type τ if $\Gamma \vdash \varphi{:}\tau$ is derivable using the rules of Fig. 2.4. An HFL-formula φ has type τ if $\emptyset \vdash \varphi{:}\tau$ is derivable using these rules. Such a type derivation is unique, provided that all lambda-abstractions and fixpoint constructions are properly typed [VV04]. It is assumed that every HFL-formula in this work is well-typed and well-named. To improve readability, the typing information is suppressed most of the time.

Definition 2.3.4. Let φ be an HFL-formula. φ is a formula of HFL^k ($\varphi \in$ HFL^k)

for some $k \in \mathbb{N}$ if there is no subformula $\psi \in \mathrm{sub}(\varphi)$ with a type of order more than k.

Common HFL-formulas tend to be hard to read, especially if they contain higher-order subformulas. The following introduces some abbreviations, making it easier to read and understand HFL-formulas:

$$\begin{aligned}
&\lambda x_1, x_2, \dots, x_l.\psi := \lambda x_1.\lambda x_2.\cdots\lambda x_l.\psi,\\
&f\, x_1\, x_2\, \cdots\, x_l := (\cdots(f\, x_1)\, x_2)\cdots)\, x_l),\\
&\circ := \lambda f_1.\lambda f_2.\lambda x.f_1\,(f_2\, x),\\
&\psi_1 \circ \psi_2 := \circ\, \psi_1\, \psi_2 \text{ and}\\
&\psi_1 \circ \psi_2 \circ \cdots \circ \psi_l := (\psi_1 \circ (\psi_2 \circ \cdots \circ (\psi_{l-1} \circ \psi_l)\cdots).
\end{aligned}$$

Purposefully, the types of the variables are suppressed as the abbreviations work for arbitrary types with exception of f_1, f_2 in the abbreviation $\circ$, which need to match each other.

The unfolding of a HFL-formula as seen for μ-calculus formulas in Def. 2.2.6 is defined similarly for HFL-formulas. Let $\varphi \in$ HFL be a formula with k distinct fixpoint subformulas. If $\varphi = \sigma x_i.\, \psi_i$ of type $\tau_0 \to \tau_1 \to \cdots \to \tau_{l-1} \to \bullet$ the unfolding for $s = (c_0, \dots, c_{k-1})$ is defined as

$$(\sigma x_i.\psi_i)^s := \begin{cases} \lambda y_0, \dots, y_{l-1}.\bot & \text{if } \sigma = \mu, c_i = 0,\\ \lambda y_0, \dots, y_{l-1}.\top & \text{if } \sigma = \nu, c_i = 0,\\ \psi_i[(\sigma x_i.\psi_i)^{s[c_i-1/c_i]}/x_i]^{s[0/c_i]} & \text{otherwise.} \end{cases}$$

The case $\varphi^s = (\psi_1\, \psi_2)^s$ is to be understood as $\psi_1^s\, \psi_2^s$. Other cases of φ are not effected by the unfolding, respectively pass the unfolding on to subformulas. With this, the finite convergence criterion of a $\mathcal{L}_\mu$-formula and LTS, formally defined in Def. 2.2.9, is defined for HFL-formulas and LTS in the same way. A noteworthy difference is that the condition on fixpoint subformulas requires that the unfolding of a fixpoint subformula defines the correct semantics of its corresponding type. For example, if a fixpoint subformula is of type $\bullet \to \bullet$ the unfolding needs to define the same function from $\bullet$ to $\bullet$ as the function of its true semantics.

2.3.3 Semantics

Given an LTS $\mathcal{T}$, an environment η is a possibly partial map which assigns to each variable of the variable set an object of its respective type. Formally, η is called Γ-respecting for some context $\Gamma = x_1^{v_1,\tau_1}, \dots, x_n^{v_n,\tau_n}$ if for every x_i holds $\eta(x_i) \in \mathcal{T}[\![\tau_i]\!]$. An update of η is denoted by $\eta[x \mapsto d]$, which maps x to d and everything else according to η.

Definition 2.3.5. Let $\mathcal{T}$ be an LTS, Γ a context and η a Γ-respecting environment. The semantics $[\![\Gamma \vdash \varphi{:}\tau]\!]_\eta^\mathcal{T}$ for an HFL-formula φ are defined inductively as given in Fig. 2.5.

If an HFL-formula φ is of type τ, it means that there is a type-derivation using the empty context. In this case its semantics regarding LTS $\mathcal{T}$ and environment η are denoted by $[\![\varphi]\!]_\eta^\mathcal{T}$. As seen in Sec. 2.3.1, each HFL-type forms a complete lattice.

$$
\begin{aligned}
[\![\Gamma \vdash p{:}\bullet]\!]_\eta^\mathcal{T} &= \{s \in S \mid p \in L(s)\} \\
[\![\Gamma \vdash x{:}\tau]\!]_\eta^\mathcal{T} &= \eta(x) \\
[\![\Gamma \vdash \neg\varphi{:}\bullet]\!]_\eta^\mathcal{T} &= S \setminus [\![\Gamma^- \vdash \varphi{:}\bullet]\!]_\eta^\mathcal{T} \\
[\![\Gamma \vdash \varphi_1 \vee \varphi_2{:}\bullet]\!]_\eta^\mathcal{T} &= [\![\Gamma \vdash \varphi_1{:}\bullet]\!]_\eta^\mathcal{T} \cup [\![\Gamma \vdash \varphi_2{:}\bullet]\!]_\eta^\mathcal{T} \\
[\![\Gamma \vdash \langle a\rangle\varphi{:}\bullet]\!]_\eta^\mathcal{T} &= \left\{s \in S \mid \text{ex. } s' \in [\![\Gamma \vdash \varphi{:}\bullet]\!]_\eta^\mathcal{T} \text{ s.t. } s \xrightarrow{a} s'\right\} \\
[\![\Gamma \vdash \lambda(x^{v,\tau}).\ \varphi{:}\tau^v \to \tau']\!]_\eta^\mathcal{T} &= f \in \mathcal{T}[\![\tau^v \to \tau']\!] \text{ s.t.} \\
&\qquad \text{f.a. } d \in \mathcal{T}[\![\tau]\!] : f(d) = [\![\Gamma, x^{v,\tau} \vdash \varphi{:}\tau']\!]_{\eta[x\mapsto d]}^\mathcal{T} \\
[\![\Gamma \vdash \varphi_1\, \varphi_2{:}\tau']\!]_\eta^\mathcal{T} &= [\![\Gamma \vdash \varphi_1{:}\tau^v \to \tau']\!]_\eta^\mathcal{T} \left([\![\Gamma^v \vdash \varphi_2{:}\tau]\!]_\eta^\mathcal{T}\right) \\
[\![\Gamma \vdash \mu(x^\tau).\ \varphi{:}\tau]\!]_\eta^\mathcal{T} &= \bigsqcap \left\{d \in \mathcal{T}[\![\tau]\!] \mid [\![\Gamma, x^+ : \tau \vdash \varphi{:}\tau]\!]_{\eta[x\mapsto d]}^\mathcal{T} \sqsubseteq_{\mathcal{T}[\![\tau]\!]} d\right\}
\end{aligned}
$$

FIGURE 2.5: The semantics of HFL.

Thus, the Knaster-Tarski Theorem [Tar+55] guarantees the existence of fixpoints for every derivable HFL fixpoint-formula.

2.3.4 The Modal μ-Calculus as a Fragment of HFL

A hierarchy for HFL-formulas which depends on the highest type occurring in its set of subformulas is given in Def. 2.3.4. An HFL-formula φ is a formula of the fragment HFL^0 if all subformulas $\psi \in \text{sub}(\varphi)$ can be typed with the base-type. Especially, if $\varphi \in \text{HFL}^0$ then there is no subformula of the form $\lambda x^{v,\tau}.\psi$ or $(\psi_1\, \psi_2)$.

Remark 2.3.6. For each $\varphi \in \text{HFL}^0$ there is an equivalent $\varphi' \in \mathcal{L}_\mu$ such that for all LTS $\mathcal{T}$ and environments η it holds that $[\![\varphi]\!]_\eta^\mathcal{T} = [\![\varphi']\!]_\eta^\mathcal{T}$, $\text{ad}(\varphi) = \text{ad}(\varphi')$ and $\varphi \in \text{FC}(\mathcal{T})$ if and only if $\varphi' \in \text{FC}(\mathcal{T})$. The same is valid for the other direction.

With some further explanation this remark becomes clear: Every $\mathcal{L}_\mu$-formula can be interpreted as a HFL-formula by using the abbreviations of Sec. 2.3.2 and assuming that every subformula is of the base-type $\bullet$.

The other direction, namely interpreting a formula $\varphi \in \text{HFL}^0$ as a $\mathcal{L}_\mu$-formula, seems to be more difficult to argue as the HFL syntax allows negations in front of arbitrary formulas. But as φ is a formula of the fragment HFL^0 it can be argued that there is no subformula of the form $\lambda x^{v,\tau}.\psi$ or $(\psi_1\, \psi_2)$. For the other cases of non-atomic negations φ can be altered according to De Morgan's laws and the duality of $[a]$ and $\langle a\rangle$ as well as μ and ν.

2.4 Alternating Parity Automata

The following definition of an alternating parity automaton is tailored to a translation between $\mathcal{L}_\mu$-formulas and alternating parity automata. A similar definition can be found in [GKL14]. Other suitable definitions can be found in [Wil01].

Let P be some set of propositions. $\mathcal{B}^+(P)$ denotes the set of positive Boolean formulas over P. The set $\neg P$ is the set of negated propositions of P, namely $\neg P = \{\neg p \mid p \in P\}$.

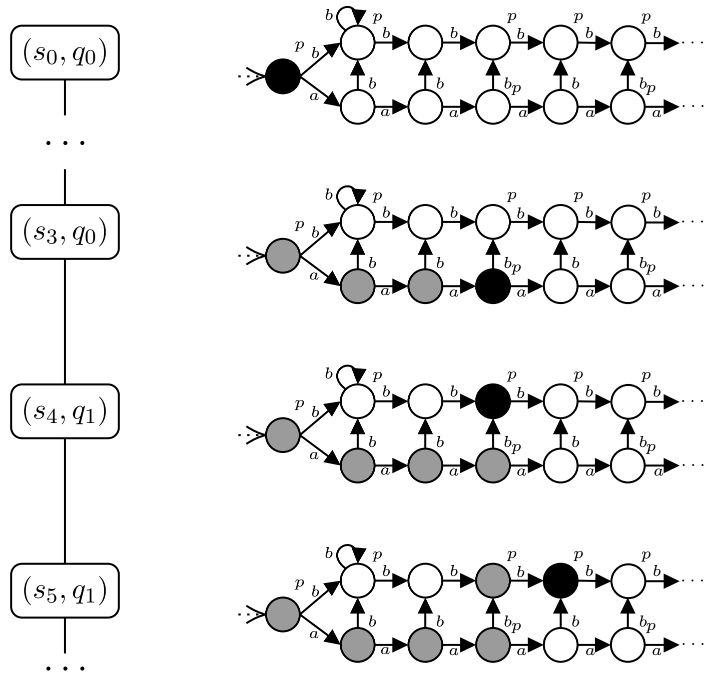

FIGURE 2.6: Visualisation of a run of automaton $\mathcal{A}$ of Ex. 2.4.2 on the LTS of Ex. 2.2.10. The black filled state is currently considered. Previous steps of the automaton on the LTS are depicted with a grey filled state.

Definition 2.4.1. Let P be a set of propositions and A a set of actions. An Alternating Parity Automaton (APA) $\mathcal{A}$ is a tuple $\mathcal{A} = (Q, q_{\text{init}}, \delta, \Omega)$, where

- Q is a finite set of states, $q_{\text{init}} \in Q$ is called the initial state,
- $\delta : Q \to \mathcal{B}^+(P \cup \neg P \cup M)$ with $M = Q \times \{\Diamond, \Box\} \times A$ is the transition function and
- $\Omega : Q \to \mathbb{N}$ the priority function.

A run ρ of $\mathcal{A}$ on some LTS $(\mathcal{T}, s_{\text{init}})$ is a tree $\rho : \mathbb{N}^* \to S \times Q$ such that $\rho(\epsilon) = (s_{\text{init}}, q_{\text{init}})$ and for each node $i \in \mathbb{N}^*$ with $\rho(i) = (s_i, q_i)$ there is a set $R \subseteq L(s_i) \cup \neg L(s_i) \cup M$ such that R satisfies $\delta(q_i)$ and it holds that:

- if $(q_j, \Diamond, a) \in R$ then there is a state $s_j \in R$ with $s_i \xrightarrow{a} s_j$ and a child $j \in \mathbb{N}^*$ of the node i such that $\rho(j) = (s_j, q_j)$,
- if $(q_j, \Box, a) \in R$ then for each $s_j \in R$ with $s_i \xrightarrow{a} s_j$ there is a child $j \in \mathbb{N}^*$ of the node i such that $\rho(j) = (s_j, q_j)$.

A run is called accepting if for every branch π it holds that the value of $\max(\{\Omega(q) \mid (s, q)$ for some s occurs infinitely often on $\pi\})$ is even. If a runs is not accepting then it is called rejecting. The language $L_{\mathcal{T}}(\mathcal{A})$ of an automaton $\mathcal{A}$ for some LTS $\mathcal{T}$ is defined as $L_{\mathcal{T}}(\mathcal{A}) = \{s \in S \mid$ there is an accepting run of $\mathcal{A}$ on $(\mathcal{T}, s)\}$.

Example 2.4.2. Let $P = \{p\}$ be a set of one proposition and $A = \{a, b\}$ a set of actions. Consider the automaton $\mathcal{A} = (\{q_0, q_1\}, q_0, \delta, \Omega)$ such that

- the transition function δ is given by $\delta(q_0) = (p \wedge (q_1, \Box, b)) \vee (q_0, \Diamond, a), \delta(q_1) = p \wedge (q_1, \Box, b)$ and
- the priority function Ω is given by $\Omega(q_1) = 0$ and $\Omega(q_0) = 1$.

A run of the automaton $\mathcal{A}$ on some LTS is shown in Fig. 2.6. In this case the run consists of a single branch. Furthermore, the run is accepting as the state q_1 is the only state that occurs infinitely often and $\Omega(q_1) = 0$ holds. There are other accepting runs, namely runs which follow the a-path a greater but finite amount of steps.

2.4.1 The Modal μ-Calculus and Alternating Parity Automata

In the following a translation from guarded $\mathcal{L}_\mu$-sentences to semantically equivalent APA is presented. Therefore, let $\mathrm{ai}_\varphi(x)$ be defined as the least natural number that satisfies

- $\mathrm{ai}_\varphi(x)$ is odd if $\mathrm{fp}_\varphi(x) = \mu x.\, \psi$ with $\psi \in \mathrm{sub}(\varphi)$,
- $\mathrm{ai}_\varphi(x)$ is even if $\mathrm{fp}_\varphi(x) = \nu x.\, \psi$ with $\psi \in \mathrm{sub}(\varphi)$ and
- for all $y \in \mathrm{sub}(\varphi)$ such that y occurs freely in $\mathrm{fp}_\varphi(x)$ holds $\mathrm{ai}_\varphi(x) \geq \mathrm{ai}_\varphi(y)$.

The following theorem can be found in a more detailed version in [GKL14]. The result is presented here again as it includes the specific translation from guarded $\mathcal{L}_\mu$-sentences to APA. This is needed in detail throughout this work.

Theorem 2.4.3. Let $\varphi \in \mathcal{L}_\mu$ be a guarded sentence, $\mathcal{T}$ an LTS and η an environment. There is an APA $\mathcal{A}_\varphi$ such that $[\![\varphi]\!]_\eta^\mathcal{T} = L_\mathcal{T}(\mathcal{A}_\varphi)$.

Proof. Let $\varphi \in \mathcal{L}_\mu$ be a guarded sentence. The components of the equivalent automaton $\mathcal{A}_\varphi = (Q, q_{\mathrm{init}}, \delta, \Omega)$ are given by

- $Q := \mathrm{sub}(\varphi) \times \{0, \ldots, \mathrm{ad}(\varphi)\}$, $q_{\mathrm{init}} = (\varphi, 0)$,
- $\Omega : Q \to \mathbb{N}$ with $\Omega(\psi, d) = d$ and
- the transition function is defined as $\delta((\psi, d)) = \mathrm{trav}_0(\psi)$, where $\mathrm{trav}_d(\psi)$ is defined as

$$\mathrm{trav}_d(\psi) = \begin{cases} \psi & \text{if } \psi = p, \neg p, \top, \bot \text{ with } p \in P, \\ \mathrm{trav}_{\mathrm{ai}_\varphi(x)}(\mathrm{fp}_\varphi(x)) & \text{if } \psi = x \text{ with } x \in V, \\ \mathrm{trav}_d(\psi') * \mathrm{trav}_d(\psi'') & \text{if } \psi = \psi' * \psi'', * \in \{\vee, \wedge\} \\ ((\psi', d), \Diamond, a) & \text{if } \psi = \langle a \rangle \psi' \text{ with } a \in A, \\ ((\psi', d), \Box, a) & \text{if } \psi = [a]\psi' \text{ with } a \in A, \\ \mathrm{trav}_d(\psi') & \text{if } \psi = \sigma x.\, \psi'. \end{cases}$$

It can be seen in the definition of $\mathrm{trav}_\psi(d)$ why guardedness is needed: If there would be an unguarded variable x the corresponding call of $\mathrm{trav}_d(x)$ would lead to an infinite chain of $\mathrm{trav}_d(x)$ calls and thus its value would be undefined. $\Box$

Example 2.4.4. Consider the formula $\varphi = \mu x.\, (\nu y.\, p \wedge [b]y) \vee \langle a \rangle x$. As seen in Ex. 2.2.3, it holds that $\text{ad}(\varphi) = 1$. The fixpoint dependency measure is given as $\text{ai}_\varphi(x) = 1$ and $\text{ai}_\varphi(y) = 0$. Note, that the condition $\text{ai}_\varphi(x) \geq \text{ai}_\varphi(y)$ is not demanded, but that $\text{ai}_\varphi(x)$ is odd and $\text{ai}_\varphi(y)$ is even. The corresponding APA $\mathcal{A}_\varphi$ is defined as $(\text{sub}(\varphi) \times \{0, 1\}, (\varphi, 0), \delta, \Omega)$, where the transition function δ is built according to trav. For example, the values of $\delta((\varphi, 0))$, $\delta((y, 0))$, $\delta(x, 0)$, $\delta(y, 1)$ and $\delta(x, 1)$ are

$$\begin{aligned}
\delta((\varphi, 0)) &= (p \wedge ((y, 0), \Box, b)) \vee ((x, 0), \Diamond, a), \\
\delta((y, 0)) &= p \wedge ((y, 0), \Box, b), \\
\delta((x, 0)) &= (p \wedge ((y, 1), \Box, b)) \vee ((x, 1), \Diamond, a), \\
\delta((y, 1)) &= p \wedge ((y, 0), \Box, b), \\
\delta((x, 1)) &= (p \wedge ((y, 1), \Box, b)) \vee ((x, 1), \Diamond, a).
\end{aligned}$$

A simplified version of the automaton $\mathcal{A}_\varphi$ is the automaton $\mathcal{A}$ of Ex. 2.4.2. The states $(y, 0)$ and $(x, 1)$ correspond to states q_1 respectively q_0. The other states mentioned above, namely $(\varphi, 0), (x, 0)$ and $(y, 1)$, have no significant influence on the language of the automaton $\mathcal{A}_\varphi$ as they are reachable at most once. These states can be seen as artifacts of the recursive definition of the transition function δ.

There is an important observation about a μ-unipolar fixpoint sentences which is used in later proofs. Let φ be a μ-unipolar sentence. In the automaton $\mathcal{A}_\varphi$ every fixpoint necessarily comes with an odd priority. This means that on a corresponding run there are only odd priorities and the priority 0, which occurs because of intermediate, modal steps. But these intermediate steps can only occur finitely often before going into another fixpoint state, which means that they do not influence the acceptance of an infinite branch. Thus, the priority of all states can be mapped to a single odd priority e.g. priority 1 without changing the language of the automaton. This leads to an alternative, simpler acceptance condition for the automaton $\mathcal{A}_\varphi$: A run of $\mathcal{A}_\varphi$ is accepting if it is finite, which in turn leads to a useful simplification of $\mathcal{A}_\varphi$ by forgetting about the priorities at all.

Observation 2.4.5. Let φ be a μ-unipolar, guarded sentence and $\mathcal{A}_\varphi = (\text{sub}(\varphi) \times \{0, \dots, \text{ad}(\varphi)\}, (\varphi, 0), \delta, \Omega)$ the corresponding APA. There is a simplified automaton $\mathcal{A}'_\varphi = (\text{sub}(\varphi), \varphi, \delta')$, where δ' is defined according to δ, but each proposition $((\psi, d), \Diamond, a)$ or $((\psi, d), \Box, a)$ is projected to $(\psi, \Diamond, a)$ respectively $(\psi, \Box, a)$. A run is defined similarly, but called accepting if it is finite. Then it holds for all LTS $\mathcal{T}$ that $L_\mathcal{T}(\mathcal{A}_\varphi) = L_\mathcal{T}(\mathcal{A}'_\varphi)$.

This simplified model is used in the proof of Sec. 4.1.

A similar observation can be made about ν-unipolar sentences. Let φ be a ν-unipolar sentence. First, notice that all existing runs of $\mathcal{A}_\varphi$ on some LTS are necessarily accepting. To ensure the existence of rejecting runs the APA model needs to be extended by a state q_{rej} which has an odd priority and its only successor is q_{rej} itself. Then, if a transition condition is not satisfiable the automaton can go into q_{rej} which does not change the language of the automaton over a given LTS. Then the observation can be made that a run is rejecting if it is finite which leads to a similar simplification of $\mathcal{A}_\varphi$ as in the μ-unipolar case. This extended and simplified model is needed for the dual case of the proof of Sec. 4.1.

2.5 Linear-Time Variants: Mu-Calculus, HFL and APA

The goal of this work is presenting examples of structures, which witness a fundamental difference in the convergence behaviour of $\mathcal{L}_\mu$- and HFL-formulas. Chap. 3 presents such examples but in the form of infinite words. With the discussion about linear-time structures and infinite words in Rem. 2.1.5 these are morally the same, but corresponding definitions of $\mathcal{L}_\mu$, HFL and APAs are needed to interpret them over infinite words.

The linear-time variant of $\mathcal{L}_\mu$ (see Sec. 2.2), called the linear-time μ-calculus $\mathcal{L}_\mu^{\text{lin}}$, is a research topic itself, originally proposed in [BKP86] and earlier mentioned in [EC80]. From this works perspective, the idea is that over linTS respectively infinite words the semantics of diamond and box coincide, so only one modality operator is needed.

Definition 2.5.1. Let P be a countable set of propositions and V a countable set of variables. A formula φ is a linear-time μ-calculus formula ($\varphi \in \mathcal{L}_\mu^{\text{lin}}$) if it is producible by the following grammar:

$$\varphi ::= q \mid \neg q \mid x \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \odot\, \varphi \mid \mu x.\, \varphi \mid \nu x.\, \varphi$$

where $q \in P$ and $x \in V$.

As mentioned above, the difference to the syntax of the usual μ-calculus lies in the modal operators. The operators $\langle a \rangle$ and $[a]$ are not present, but a general one $\odot$ is.

With regard to the introduction of this section, the semantics of $\mathcal{L}_\mu^{\text{lin}}$ are defined in context of infinite words.

Definition 2.5.2. Let $w \in \Sigma^w$ with $w = x_0\, x_1\, \ldots$ be an infinite word, $\eta : V \to 2^{\mathbb{N}}$ an environment and φ a $\mathcal{L}_\mu^{\text{lin}}$-formula. The semantics of φ over w given η are denoted by $[\![\varphi]\!]_\eta^w$ and inductively defined as

$$\begin{aligned}
&[\![a]\!]_\eta^w := \{i \in \mathbb{N} \mid x_i = a\}, \quad [\![\neg a]\!]_\eta^w := \{i \in \mathbb{N} \mid x_i \neq a\},\\
&[\![x]\!]_\eta^w := \eta(x), \quad [\![\bot]\!]_\eta^w = \emptyset, \quad [\![\top]\!]_\eta^w = \mathbb{N},\\
&[\![\psi_1 \vee \psi_2]\!]_\eta^w := [\![\psi_1]\!]_\eta^w \cup [\![\psi_2]\!]_\eta^w, \quad [\![\psi_1 \wedge \psi_2]\!]_\eta^w := [\![\psi_1]\!]_\eta^w \cap [\![\psi_2]\!]_\eta^w,\\
&[\![\odot\, \varphi]\!]_\eta^w := \{i \in \mathbb{N} \mid i+1 \in [\![\varphi]\!]_\eta^w\},\\
&[\![\mu x.\, \varphi]\!]_\eta^w := \bigcap \{U \subseteq \mathbb{N} \mid [\![\varphi]\!]_{\eta[x \mapsto U]}^w \subseteq U\},\\
&[\![\nu x.\, \varphi]\!]_\eta^w := \bigcup \{U \subseteq \mathbb{N} \mid U \subseteq [\![\varphi]\!]_{\eta[x \mapsto U]}^w\}.
\end{aligned}$$

Further definitions and abbreviations regarding syntax of $\mathcal{L}_\mu$-formulas (see Sec. 2.2.1), the alternation hierarchy given for $\mathcal{L}_\mu$-formulas (see Sec. 2.2.2) and the guardedness procedure (see Sec. 2.2.5) can be defined in the same manner for $\mathcal{L}_\mu^{\text{lin}}$. The same holds for the definition of unfolding a formula and the finite convergence criterion (see Sec. 2.2.4). The latter is then denoted by FC(w) for some word $w \in \sigma^\omega$.

It is noteworthy that for all $\varphi \in \mathcal{L}_\mu^{\text{lin}}$ and infinite words $w \in \Sigma^\omega$ it holds that $[\![\varphi]\!]_\eta^w = [\![\varphi, \omega]\!]_\eta^w$, which means that the ω-approximation of φ is equal to its true semantics. The reason for this is that w is an infinite word which means that every position respectively state has finitely many successors. For finitely-branching systems each fixpoint iteration terminates at the latest with the ω-step.

Remark 2.5.3. For every $\mathcal{L}_\mu^{\text{lin}}$-formula and infinite word w there is a formula $\varphi' \in \mathcal{L}_\mu$ (and vice-versa) such that for all environments η it holds that: $[\![\varphi]\!]_\eta^w = [\![\varphi']\!]_\eta^{\mathcal{K}_w}$, $\text{ad}(\varphi) = \text{ad}(\varphi')$ and $\varphi \in \text{FC}(w)$ if and only if $\varphi' \in \text{FC}(\mathcal{K}_w)$.

A linear-time variant of HFL (see Sec. 2.3) called HFL_{lin} can be defined in the exact same way. To avoid to much repetition it is referred to the relation between $\mathcal{L}_\mu$ and $\mathcal{L}_\mu^{\text{lin}}$. The relation between HFL and HFL_{lin} is the same, which means that the operators $\langle a\rangle$ and $[a]$ coincide to $\odot$ and the semantics of HFL_{lin} are defined over infinite words.

Remark 2.5.4. For every HFL_{lin}-formula and infinite word w there is a formula $\varphi' \in$ HFL (and vice-versa) such that for all environments η it holds that: $[\![\varphi]\!]_\eta^w = [\![\varphi']\!]_\eta^{\mathcal{K}_w}$, $\text{ord}(\varphi) = \text{ord}(\varphi')$ and $\varphi \in \text{FC}(w)$ if and only if $\varphi' \in \text{FC}(\mathcal{K}_w)$.

An alternating parity automata (see Sec. 2.4) $\mathcal{A}$ can also be interpreted over infinite words. Let $w \in \Sigma^\omega$ with $w = x_0\, x_1\, \ldots$. The transition function δ is given by $\delta : Q \to \mathcal{B}^+(\Sigma \cup \neg\Sigma \cup M)$ with $M = Q \times \{\odot\}$. A run on w is a tree $\rho : \mathbb{N}^* \to \mathbb{N} \times Q$ with root $\rho(\epsilon) = (0, q_{\text{init}})$. The conditions on children-nodes are equal to the LTS-case but with $R := \{x_i\} \cup \neg\{x_i\} \cup M$. The acceptance condition for a run stays the same and the language of $\mathcal{A}$ for a given word w is defined as $L_w(\mathcal{A}) := \{v \mid v \in \text{suff}(w) \text{ and there is an accepting run of } \mathcal{A} \text{ on } v\}$. With this direct relation to Def. 2.4.1 it immediately follows from the translation given in Th. 2.4.3 that for each $\mathcal{L}_\mu^{\text{lin}}$-formula there is a semantically equivalent APA.

Corollary 2.5.5. Let $\varphi \in \mathcal{L}_\mu^{\text{lin}}$, $w \in \Sigma^\omega$ and η an environment. There is an APA $\mathcal{A}_\varphi$ such that it holds $[\![\varphi]\!]_\eta^w = L_w(\mathcal{A})$.

There is an important implication of these relations for this work. All differences in the finite convergence of $\mathcal{L}_\mu^{\text{lin}}$-formulas and HFL_{lin}-formulas over some family of infinite words imply that there is a class of linear-time structures that witness the same difference between $\mathcal{L}_\mu$-formulas and HFL-formulas. Thus, further considerations can be restricted to these linear-time variants.

Chapter 3

Words with Finite Convergence of μ-Formulas

This chapter introduces a family of words $w_{n,m}$ which witness a fundamental difference in the convergence behaviour of the modal μ-calculus and the higher-order fixpoint logic. It is strongly suspected that all $\mathcal{L}_{\mu}^{\text{lin}}$-formulas are finitely converging over $w_{n,m}$ (see Chap. 4). In contrast to this there are HFL_{lin}-formulas which are not finitely converging for some $w_{n,m}$ (see Chap. 5). The implications of this difference in the convergence behaviour are discussed in Chap. 6

Definition 3.0.1. Let $\Sigma = \{a, b\}$ be a binary alphabet. An infinite word $w_{n,m} \in \Sigma^\omega$ is defined as $w_{n,m} = \alpha_0\, \alpha_1 \cdots$ with

$$\alpha_0 = a, \qquad \alpha_{i+1} = \alpha_i^n \beta_i^m \alpha_i^n,$$
$$\beta_0 = b, \qquad \beta_{i+1} = \beta_i^n \alpha_i^m \beta_i^n,$$

where $n, m \in \mathbb{N}$ such that $n > m$, $m \geq 1$.

In the remaining part of this work let $w_{n,m}$ refer to an instance of this family. To avoid trivial instances of this family, m is restricted to be at least 1. The restriction $n > m$ is a necessary feature. It ensures that the middle part of some α_i, namely β_{i-1}^m, is a prefix of β_{i+1}. In detail this property is needed in Lem. 4.1.7.

Example 3.0.2. Let $n = 2$ and $m = 1$. The respective word $w_{2,1}$ is given by

$$w_{2,1} = \underbrace{a}_{\alpha_0}\underbrace{\underbrace{a}_{\alpha_0}\underbrace{a}_{\alpha_0}\underbrace{b}_{\beta_0}\underbrace{a}_{\alpha_0}\underbrace{a}_{\alpha_0}}_{\alpha_1}\underbrace{\underbrace{a\,a\,b\,a\,a}_{\alpha_1}\underbrace{a\,a\,b\,a\,a}_{\alpha_1}\underbrace{b\,b\,a\,b\,b}_{\beta_1}\underbrace{a\,a\,b\,a\,a}_{\alpha_1}\underbrace{a\,a\,b\,a\,a}_{\alpha_1}}_{\alpha_2}\cdots.$$

The following proposition describes important properties of a word $w_{n,m}$. For each quantitative property a sufficiently tight but not optimal bound is given.

Proposition 3.0.3. Let $w_{n,m}$ be a word of Def. 3.0.1 and $i \in \mathbb{N}$. It holds that

1. the length of α_i is $|\alpha_i| = |\beta_i| = (2n+m)^i$,
2. the length of $u \in \text{pre}(w_{n,m})$ with $u = \alpha_0\, \alpha_1 \cdots \alpha_{i-1}$ is $|u| = \frac{1-(2n+m)^i}{1-(2n+m)} \leq (2n+m)^i$,
3. there is no $u \in \text{suff}(w_{n,m})$ such that $\alpha_i \notin \inf(u)$,
4. the length of $u \in \inf(w_{n,m})$ with $\alpha_i \notin \inf(u)$ is $|u| \leq (2+m)(2n+m)^i$ and

5. the length of all $u \in \inf(w_{n,m})$ with $u = \gamma_0\,\gamma_1\,\cdots\,\gamma_{k-1}$ such that $\gamma_j = \alpha_i$ for all $j < k$ is $|u| = 2n(2n+m)^i$.

Proof. First, by induction it is proven for all $i \in \mathbb{N}$ that $|\alpha_i| = |\beta_i| = (2n+m)^i$. The case $i = 0$ is trivial. Let $i = j+1$. Per definition it follows that $|\alpha_i| = 2n|\alpha_j| + m|\beta_j| = 2n(2n+m)^j + m(2n+m)^j = (2n+m)^{j+1}$. The case β_i is shown in the exact same way.

This result can be used to prove the second property. From the first property follows that $|u| = \sum_{j=0}^{i-1}(2n+m)^j$. This is a partial sum of the geometric series with base $(2n+m)$. Thus, it follows that $\sum_{j=0}^{i-1}(2n+m)^j = \frac{1-(2n+m)^i}{1-(2n+m)} \leq (2n+m)^i$. The validity of the inequality is inferred in the following way:

$$\begin{aligned} & \frac{1-(2n+m)^i}{1-(2n+m)} \leq (2n+m)^i \\ \Leftrightarrow\quad & (2n+m)^i \leq (2n+m)^{i+1} - (2n+m)^i \\ \Leftrightarrow\quad & 1 \leq 2n+m-1, \end{aligned}$$

which is true as $n, m > 0$.

The third property is a straightforward implication of Def. 3.0.1.

To prove the fourth property, consider an infix $u \in \inf(w_{n,m})$. It follows that u is also an infix of a sequence $\alpha_j\,\alpha_{j+1}\,\cdots\,\alpha_{i+k}$ for some $j, k \in \mathbb{N}$. From Def. 3.0.1 it follows that $\alpha_i \in \text{pre}(\alpha_j)$ and $\alpha_i \in \text{suff}(\alpha_j)$ for all $j \geq i$, which means that no such α_j can be an infix of u. Thus, the maximum case is either the prefix of $w_{n,m}$ described in the second property or it holds that $|u| = (2+m)(2n+m)^i - 2$ which describes the case that u is almost equal to $\alpha_i\beta_i^m\alpha_i$. This is the biggest infix of α_{i+1} that does not completely include α_i. As the second one is greater, it gives the desired bound.

For the fifth property it is argued that there are at most $2n$ consecutive α_i in $w_{n,m}$. By induction it can be shown for all α_j with $j \geq i$ that the longest prefix and suffix of consecutive α_i in α_j is equal to n and that the longest infix is equal to $2n$. From the definition of $w_{n,m}$ it follows that u must be an infix of $\alpha_i\,\alpha_{i+1}\,\cdots\,\alpha_{i+k}$ for some $k \in \mathbb{N}$. But with the previous argument it follows that for all k there are at most $2n$ consecutive α_i in a sequence $\alpha_i\,\alpha_{i+1}\,\cdots\,\alpha_{i+k}$. This implies that the same is valid for u. □

The properties four and five are core properties of a word $w_{n,m}$ and the main reason why this family of words is defined this way. Informally put, property four states for all i that there is a bound for the distance between two occurrences of α_i in $w_{n,m}$. Property five states for all i that there is a bound for the number of consecutive α_i in $w_{n,m}$. This leads to the conjecture that all $\mathcal{L}_\mu^{\text{lin}}$-formulas are finitely converging over all $w_{n,m}$. Intuitively, the semantics of simple $\mathcal{L}_\mu$-formulas and especially of formulas with fixpoint operators can be understood as "somewhere pattern ABC is valid" (least fixpoints) or "from here on pattern DEF is always valid" (greatest fixpoints). Nested fixpoint formulas are naturally difficult to describe by such a simple intuition, but this basic intuition is transmitted to these formulas. Property four and five of a word $w_{n,m}$ ensure, regardless of the position in $w_{n,m}$, that there is a bound until "somewhere pattern ABC" and "always pattern DEF" is verified. This leads to the assumption that every fixpoint $\mathcal{L}_\mu^{\text{lin}}$-formula only needs finitely many unfoldings to reach its true semantics over $w_{n,m}$. In a following step this leads to the assumption that finite convergence is guaranteed for all $\mathcal{L}_\mu^{\text{lin}}$-formulas over all $w_{n,m}$. As mentioned above this is investigated in Chap. 4. In contrast to this, fixpoint formulas of higher-order fragments of HFL_{lin} are able to modify the pattern that has to be valid "somewhere" or "always" from iteration step to iteration step. This leads to the assumption that there are HFL_{lin}-formulas which are not finitely converging over some $w_{n,m}$ as there are infinitely many patterns to look for in each $w_{n,m}$. This is investigated in Chap. 5.

Another important property of a word $w_{n,m}$ is that for all i any suffix from and after the first α_i in $w_{n,m}$ can be represented as a sequence of α_i and β_i. One can see this as different levels of detail of observation. This is formalised in the following proposition.

Proposition 3.0.4. Let $w_{n,m}$ be a word of Def. 3.0.1, $i \in \mathbb{N}$, $\gamma_j \in \{\alpha_i, \beta_i\}$ for all $j \in \mathbb{N}$ and $n_i = (1-(2n+m)^i)/(1-(2n+m))$ the starting point of the first α_i in $w_{n,m}$. There is a sequence $\alpha_i\, \gamma_0\, \gamma_1\, \cdots$ such that $w_{n,m}[n_i :] = \alpha_i\, \gamma_0\, \gamma_1\, \cdots$.

Proof. By induction it is shown that for all $k \geq 1$ holds that $\alpha_{i+k} = \gamma'_1\, \gamma'_2\, \cdots\, \gamma'_{(2n+m)^k}$ with $\gamma'_j \in \{\alpha_i, \beta_i\}$ for all $1 \leq j \leq (2n+m)^k$. The case $k = 1$ is a direct implication of Def. 3.0.1.

Assume that the statement holds for k. For the induction step consider

$$\alpha_{i+(k+1)} = \alpha_{i+k}\, \cdots\, \alpha_{i+k}\, \beta_{i+k}\, \cdots\, \beta_{i+k}\, \alpha_{i+k}\, \cdots\, \alpha_{i+k},$$

where α_{i+k} occurs n times, respectively β_{i+k} occurs m times in sequence. By the induction hypothesis it follows that each $\alpha_{i+k}, \beta_{i+k}$ can be represented as a sequence of $(2n+m)^k$ blocks of α_i and β_i. But this means that $\alpha_{i+(k+1)}$ can be represented by $(2n+m) \cdot (2n+m)^k = (2n+m)^{k+1}$ blocks of α_i and β_i.

With this statement proven it follows that every α_{i+k} with $k \geq 1$ can be represented as a sequence of α_i and β_i. This implies the existence of a sequence $\gamma_0\, \gamma_1\, \cdots$ such that $w_{n,m}[n_i :] = \alpha_i\, \gamma_0\, \gamma_1\, \cdots$. □

Chapter 4

Finite Convergence of μ-Calculus Formulas

The goal of this chapter is to show that every $\mathcal{L}_\mu^{\text{lin}}$-sentence is finitely converging over each $w_{n,m}$. A formal proof is given for all alternation-free sentences. An idea of a proof for the general case is sketched.

The first step (Sec. 4.1) is to prove the finite convergence of all unipolar $\mathcal{L}_\mu^{\text{lin}}$-sentences over all $w_{n,m}$. Let φ be a unipolar sentence. The proof is divided into the following parts: First, the guarded variant φ' of φ is considered. Then, a relation between the height of a run of $\mathcal{A}_{\varphi'}$ and a finite unfolding of φ' is established, allowing to proceed the proof by arguing on $\mathcal{A}_{\varphi'}$. After this, the automaton $\mathcal{A}_{\varphi'}$ is translated into a deterministic automaton $\mathcal{D}_{\varphi'}$ by making an intermediate translation into a nondeterministic automaton $\mathcal{N}_{\varphi'}$. A relation between the height and length of corresponding runs is established in each translation. After this it is argued for all $w_{n,m}$ and all $v \in \text{suff}(w_{n,m})$ that there is a bound for an accepting run of $\mathcal{D}_{\varphi'}$ on v. Applying all translation in the opposite direction results in the fact that for all $w_{n,m}$ there is a unfolding φ'^k which is semantically equivalent to φ' over $w_{n,m}$. In the following this is used to prove that φ' is finitely converging over all $w_{n,m}$. In the end it is argued that finite convergence of the guarded variant φ' implies finite convergence of the original, possibly unguarded sentence φ. In conclusion, this proves for all unipolar $\mathcal{L}_\mu^{\text{lin}}$-sentences that they are finitely converging over all $w_{n,m}$.

The second step (Sec. 4.2) is to prove the finite convergence of all alternation-free $\mathcal{L}_\mu^{\text{lin}}$-sentences over all $w_{n,m}$. Let φ be an alternation-free sentence. First, it can be observed that φ is a stack of unipolar sentences. Then, finite convergence of φ over all $w_{n,m}$ is shown by induction while using the result that all unipolar sentences are finitely converging over all $w_{n,m}$.

The third and last step (Sec. 4.3) is to show finite convergence of all $\mathcal{L}_\mu^{\text{lin}}$-sentences over all $w_{n,m}$. As a formal proof is out of scope for this work, this step is only sketched. Let φ be an arbitrary sentence. It is a known that for all infinite words there is an alternation-free sentence φ' which is semantically equivalent to φ. Applying the previous results of this work on φ' results in the fact that for each $w_{n,m}$ there is a sentence of basic modal logic $\varphi'_{n,m}$ which is semantically equivalent to φ. Then, a bound d for the length of words which are included in the semantics of $\varphi'_{n,m}$ can be derived. This bound is also valid for the semantics of φ over $w_{n,m}$. This means that φ can only distinguish between infixes of $w_{n,m}$ up to length d. With this, it can be argue that there are semantically equivalent unfoldings of all outer fixpoint subformulas in φ, while assuming that inner fixpoint subformulas are computed completely. This argument could be used in an inductive approach, but this comes with some difficulties which are not resolved in this work.

4.1 Finite Convergence of Unipolar Fixpoint Formulas

The technically most demanding proof of this work is to argue that all unipolar $\mathcal{L}_\mu^{\text{lin}}$-sentences are finitely converging over all $w_{n,m}$. The first step in this proof is to show for al guarded, unipolar sentences $\varphi \in \mathcal{L}_\mu^{\text{lin}}$ and words $w_{n,m}$ that there is a semantically equivalent unfolding of φ. Thus, let φ be guarded and unipolar. For clarity reasons the proof is done in detail for the μ-unipolar case. With a few technical adaptions the proof for the ν-case is dual to the following. This is discussed in detail at the end of this section.

First, it is shown for all $w_{n,m}$ that there is a relation between the height of an accepting run of $\mathcal{A}_\varphi$ on any suffix of $w_{n,m}$ and the semantics of some unfolding of φ. This allows proceeding the proof with automata-theoretic arguments. To establish this relation, the simplified but equivalent version of $\mathcal{A}_\varphi$ is used, which is presented in Obs. 2.4.5. This simplification forgets about priorities and calls a run accepting if it is finite, which is valid as φ is $\mathcal{L}_\mu^{\text{lin}}$-unipolar. Therefore, let $\mathcal{A}_\varphi = (\text{sub}(\varphi), \varphi, \delta)$ with $\delta : \text{sub}(\varphi) \to \mathcal{B}^+(\Sigma \cup \neg\Sigma \cup (\text{sub}(\varphi) \times \{\odot\}))$. Furthermore, the semantics of $\mathcal{A}_\varphi$ are considered in context of infinite words as described in Sec. 2.5.

Lemma 4.1.1. Let $\varphi \in \mathcal{L}_\mu^{\text{lin}}$ be a guarded, μ-unipolar sentence, $w \in \Sigma^\omega$, $i \in \mathbb{N}$ and η an environment. If there is an accepting run ρ of $\mathcal{A}_\varphi$ on $w[i:]$ then it follows that $i \in [\![\varphi^{\text{height}(\rho)+1}]\!]_\eta^w$.

Proof. W.l.o.g. it is assumed that $\varphi = \mu x_0.\, \psi_0$ and that there are n' distinct fixpoint subformulas. First, a stronger statement is proven: If ρ is an accepting run on $w[i:]$ such that on each branch π there are at most c_j nodes labeled with (i', x_j) for some $i' \in \mathbb{N}$ it follows that $i \in [\![(\mu x_0.\, \psi_0)^{s+1}]\!]_\eta^w$ with $s = (c_0, \ldots, c_{n'-1})$. The proof is done by induction on the nesting-depth of $\mu x_0.\, \psi_0$. In the following, the variables $i', i'', \ldots$ refer to arbitrary positions of $w_{n,m}$.

Assume that $\text{nd}(\mu x_0, \psi_0) = 1$, which means that $s = (c_0)$. Let ρ be an accepting run of $\mathcal{A}_{\mu x_0.\, \psi_0}$ on $w[i:]$. If $c_0 = 0$ then there is no node labeled with (i', x_0), which means that ρ is also an accepting run of the automaton $\mathcal{A}_{\psi_0[\bot/x_0]}$ on $w[i:]$. According to Th. 2.4.3 it follows that $i \in [\![\psi_0[\bot/x_0]]\!]_\eta^w$ which is the same as $[\![(\mu x_0.\, \psi_0)^{(1)}]\!]_\eta^w$. Assume that $c_0 = k+1$ and that the stronger statement holds for k. For each node $h \in \mathbb{N}^*$ which is labeled with (i', x_0) it can be inferred that there is an accepting run of $\mathcal{A}_{\mu x_0.\, \psi_0}$ on $w[i':]$ with no more than k states labeled with (i'', x_0) on each branch. Namely, the run with root $(i', \mu x_0.\, \psi_0)$ followed by the subtree of ρ starting in node h. Thus, it follows that $i' \in [\![(\mu x_0.\, \psi_0)^{(k+1)}]\!]_\eta^w$. This means that ρ is also an accepting run of $\mathcal{A}_{\psi_0[(\mu x_0.\, \psi_0)^{(k+1)}/x_0]}$ on $w[i:]$ which leads to

$$i \in [\![\psi_0[(\mu x_0.\, \psi_0)^{(k+1)}/x_0]]\!]_\eta^w = [\![(\mu x_0.\, \psi_0)^{(k+2)}]\!]_\eta^w.$$

Assume that $\text{nd}(\mu x_0, \psi_0) = m' + 1$ and that the stronger statement holds for formulas of nesting-depth up to m'. Let ρ be an accepting run of $\mathcal{A}_{\mu x_0.\, \psi_0}$ on $w[i:]$ such that for all $0 \leq j < n'$ there are c_j occurrences of (i', x_j). This case is shown by induction on c_0. If $c_0 = 0$ then it follows that ρ is also an accepting run of $\mathcal{A}_{\psi_0[\bot/x_0]}$ on $w[i:]$. Then, the stronger statement follows by induction on the nesting-depth. Assume that $c_0 = k+1$ and that the stronger statement holds for $c_0 = k$. By induction on c_0 it follows for any node labeled with (i', x_0) that $i' \in [\![(\mu x_0.\, \psi_0)^{s[k/c_0]}]\!]_\eta^w$ with $s = (c_0, \ldots, c_{n'-1})$. This means that ρ is also an accepting run for the automaton $\mathcal{A}_{\psi[(\mu x_0.\, \psi_0)^{s[k/c_0]}/x_0]}$. Again, the statement follows by induction on the nesting-depth.

With this stronger statement proven, the statement of this lemma can be inferred from the observation that for a finite run ρ it must hold for each c_j that $c_j \leq \text{height}(\rho)$ and the monotonicity properties of Lem. 2.2.8. □

For a guarded, μ-unipolar sentence φ this lemma enables arguing on $\mathcal{A}_\varphi$, respectively its runs, to find a semantically equivalent unfolding of φ over some $w_{n,m}$. Namely, if there is $k \in \mathbb{N}$ for all $i \in \mathbb{N}$ with an accepting run ρ of $\mathcal{A}_\varphi$ on $w_{n,m}[i:]$ such that $\text{height}(\rho) \leq k$ then it holds that $[\![\varphi^k]\!]_\eta^{w_{n,m}} = [\![\varphi]\!]_\eta^{w_{n,m}}$. To find such k it is necessary to transform $\mathcal{A}_\varphi$ into a semantically equivalent, deterministic automaton.

From Alternating to Deterministic Automata

First, $\mathcal{A}_\varphi$ is translated into a semantically equivalent, pure nondeterministic automaton called $\mathcal{N}_\varphi$. The main difference between $\mathcal{A}_\varphi$ and $\mathcal{N}_\varphi$ is that runs of $\mathcal{N}_\varphi$ are no longer trees but paths. Informally put, this means that the translation from $\mathcal{A}_\varphi$ to $\mathcal{N}_\varphi$ avoids alternation. To ease later steps, $\mathcal{N}_\varphi$ is defined in the manner of a word automaton. This means that for a given state and proposition the transition function of $\mathcal{N}_\varphi$ maps into a set of possible successor states. This translation is similar to the language equivalence preserving translation from AFA to NFA. For an example of this translation see [HL11].

Definition 4.1.2. Let $A_\varphi = (\text{sub}(\varphi), \varphi, \delta)$ be the simplified APA of a guarded, μ-unipolar sentence $\varphi \in \mathcal{L}_\mu^{\text{lin}}$ with the transition function $\delta : \text{sub}(\varphi) \to \mathcal{B}^+(\Sigma \cup \neg\Sigma \cup (\text{sub}(\varphi) \times \{\odot\}))$. The automaton $\mathcal{N}_\varphi$ is defined as the tuple $(Q, q_{\text{init}}, q_{\text{acc}}, \delta')$ such that

- $Q := 2^{\text{sub}(\varphi)}$ is the state set,
- $q_{\text{init}} := \{\varphi\}$ is the initial state, $q_{\text{acc}} := \emptyset$ is called the accepting state and
- the transition function $\delta' : Q \times \Sigma \to 2^Q$ is defined as

$$\delta'(q, a) := \{q' \mid \{a\} \cup \neg\{a\} \cup q' \times \{\odot\} \text{ satisfies } \bigwedge_{\psi \in q} \delta(\psi)\}.$$

A run of $\mathcal{N}_\varphi$ on some infinite word $w \in \Sigma^\omega$ is a possibly infinite sequence $\pi = q_0\, q_1 \cdots$ such that $q_0 = q_{\text{init}}$ and for q_{j+1} with $j \geq 0$ holds $q_{j+1} \in \delta'(q_j, w(j))$. If there is $i \in \mathbb{N}$ such that $q_i = q_{\text{acc}}$ then $\pi = q_0 \cdots q_i$. Thus, a run ends with the first occurrence of q_{acc}. A run is called accepting if it reaches q_{acc} respectively if it is finite.

The following lemma proves that $\mathcal{A}_\varphi$ and $\mathcal{N}_\varphi$ are equivalent over all infinite words and that the length of an accepting run of $\mathcal{N}_\varphi$ and the height of an accepting run of $\mathcal{A}_\varphi$ do correspond.

Lemma 4.1.3. Let $w \in \Sigma^\omega$ and $i \in \mathbb{N}$. It holds that $L_w(\mathcal{A}_\varphi) = L_w(\mathcal{N}_\varphi)$ and if there is an accepting run π of $\mathcal{N}_\varphi$ on $w[i:]$ then there is an accepting run ρ of $\mathcal{A}_\varphi$ on $w[i:]$ with $\text{height}(\rho) = |\pi| - 1$.

Proof. First, language equivalence of $\mathcal{A}_\varphi$ and $\mathcal{N}_\varphi$ over w is proven. The directions $\subseteq$ and $\supseteq$ are proven separately.

Assume that $i \in L_w(\mathcal{A}_\varphi)$. Then it holds that there is an accepting run ρ of $\mathcal{A}_\varphi$ on $w[i:]$. An accepting run for the automaton $\mathcal{N}_\varphi$ can be built by assembling all states levelwise into a a sequence $q_0 q_1 \cdots$. From the definition of the transition function it follows that this is a run of $\mathcal{N}_\varphi$ and from the fact that there is maximum level k of ρ it follows that it is an accepting run. Thus, all state conditions are satisfied without using a proposition from $Q \times \{\odot\}$ in the k-th level of ρ. This implies that $q_{k+1} = \emptyset$ which is the same as q_{acc}.

Assume that $i \in L_w(\mathcal{N}_\varphi)$. Then it follows that there is a run $q_0\, q_1 \cdots q_{k-1}$ of $\mathcal{N}_\varphi$ on $w[i:]$ with $q_{k-1} = q_{\text{acc}}$. A run ρ for $\mathcal{A}_\varphi$ can be built by labeling the root with (i_0, φ) and then taking children from q_{j+1} in a step-wise fashion to build the $(j+1)$-th level of ρ according to the conjunction $\bigwedge_{\psi_j \in q_j} \delta(\psi_j)$. From the fact that $q_{k-1} = \emptyset$ it follows that the maximum level of ρ is $k-2$. This makes ρ an accepting run.

The relation between the length of an accepting run π of $\mathcal{N}_\varphi$ and corresponding run ρ of $\mathcal{A}_\varphi$ is a straightforward implication of the previous arguments. The difference of 1 is explained by the fact that the root level of a tree is 0. □

So far it has been shown that there is a nondeterministic automaton $\mathcal{N}_\varphi$ which is semantically equivalent to $\mathcal{A}_\varphi$ and that there is a relation between the length of an accepting run of $\mathcal{N}_\varphi$ and height of an accepting run of $\mathcal{A}_\varphi$. The next step is to translate $\mathcal{N}_\varphi$ into an equivalent determinstic automaton which is called $\mathcal{D}_\varphi$.

Definition 4.1.4. Let $\mathcal{N}_\varphi = (Q, q_{\text{init}}, q_{\text{acc}}, \delta)$. The deterministic automaton $\mathcal{D}_\varphi$ is a tuple $(Q', q'_{\text{init}}, q'_{\text{acc}}, \delta')$ with

- the state set $Q' := 2^Q$,
- $q'_{\text{init}} := \{q_{\text{init}}\}$, $q'_{\text{acc}} := \{q_{\text{acc}}\}$ and
- the transition function $\delta' : Q' \times 2^\Sigma \to Q'$ such that

$$\delta'(q, a) := \begin{cases} q'_{\text{acc}} & \text{if } q_{\text{acc}} \in \bigcup_{q' \in q} \delta(q', a), \\ \bigcup_{q' \in q} \delta(q', a) & \text{otherwise.} \end{cases}$$

The definition of a run and its acceptance condition are the same as for $\mathcal{N}_\varphi$.

In the same manner as in the translation from $\mathcal{A}_\varphi$ to $\mathcal{N}_\varphi$, language equivalence is given between $\mathcal{N}_\varphi$ and $\mathcal{D}_\varphi$. Furthermore, the lengths of corresponding, accepting runs are identical.

Lemma 4.1.5. Let $w \in \Sigma^\omega$ and $i \in \mathbb{N}$. It holds that $L_w(\mathcal{N}_\varphi) = L_w(\mathcal{D}_\varphi)$ and if π is an accepting run of $\mathcal{D}_\varphi$ on $w[i:]$ then there is an accepting run π' of $\mathcal{N}_\varphi$ on $w[i:]$ such that $|\pi'| = |\pi|$.

Proof. The proof is similar to the usual powerset construction on NFA. This includes language equivalence as well as a straightforward argument for the length equivalence of runs. □

The given translation from $\mathcal{A}_\varphi$ to $\mathcal{D}_\varphi$ allows proceeding the proof with $\mathcal{D}_\varphi$. The relations between the automata $\mathcal{A}_\varphi, \mathcal{N}_\varphi$ and $\mathcal{D}_\varphi$ transfer a bound for the length of accepting runs of $\mathcal{D}_\varphi$ on suffixes of some $w_{n,m}$ to the height of accepting runs of $\mathcal{A}_\varphi$ on suffixes of $w_{n,m}$. Furthermore, language equivalence is preserved in each translation. With the help of Lem. 4.1.1, this implies the existence of an unfolding φ^k which is semantically equivalent to φ over $w_{n,m}$.

Bounded Reachability of Deterministic Automata

The next step is to argue that for all $w_{n,m}$ and all $i \in L_{w_{n,m}}(\mathcal{D}_\varphi)$ the length of the accepting run of $\mathcal{D}_\varphi$ on $w_{n,m}[i:]$ is bounded.

Let $\mathcal{D}_\varphi = (Q, q_{\text{init}}, q_{\text{acc}}, \delta)$. For a state $q \in Q$ and finite word $u \in \Sigma^*$ with $u = x_0 x_1 \cdots x_{i-1}$ let $\text{reach}_q(u) := \{q' \mid \text{there is } q_0 q_1 \cdots q_i \text{ s.t. } \delta(q_j, x_j) = q_{j+1} \text{ f.a. } j < i, q_0 = q' \text{ and } q_i = q\}$. Informally, $\text{reach}_q(u)$ is a subset of Q which includes all states that reach q by walking along the path, which is defined by u, in $\mathcal{D}_\varphi$. For the state q_{acc} the reachability is defined slightly different: $\text{reach}_{q_{\text{acc}}}(u)$ is defined as the set $\{q' \mid \text{there is } q_0 q_1 \cdots q_{i'}, i' \leq i \text{ s.t } \delta(q_j, x_j) = q_{j+1} \text{ f.a. } j < i', q_0 = q' \text{ and } q_{i'} = q_{\text{acc}}\}$. This definition respects that a run of $\mathcal{D}_\varphi$ ends if it reaches q_{acc}.

First, some preliminary considerations are needed. These concern the reachability of q_{acc} in context of some $w_{n,m}$.

Lemma 4.1.6. Let $w_{n,m}$ be a word of Def. 3.0.1.

1. For all $i \in \mathbb{N}$ it holds that $\text{reach}_{q_{\text{acc}}}(\alpha_i) \subseteq \text{reach}_{q_{\text{acc}}}(\alpha_{i+1})$ and $\text{reach}_{q_{\text{acc}}}(\beta_i) \subseteq \text{reach}_{q_{\text{acc}}}(\beta_{i+1})$ and
2. there is $h \in \mathbb{N}$ such that for all $i \in \mathbb{N}$ with $i \geq h$ it holds that $\text{reach}_{q_{\text{acc}}}(\alpha_h) = \text{reach}_{q_{\text{acc}}}(\alpha_i)$ and $\text{reach}_{q_{\text{acc}}}(\beta_h) = \text{reach}_{q_{\text{acc}}}(\beta_i)$.

Proof. The first oberservation results from the structure of w as it holds for all $i \in \mathbb{N}$ that $\alpha_i \in \text{pre}(\alpha_{i+1})$. The same holds for β_i and β_{i+1}. The second observation follows from the first one and the fact that Q is finite. □

Let $(\text{reach}_{q_{\text{acc}}}(\alpha_i) \cup \text{reach}_{q_{\text{acc}}}(\beta_i)) \cup M_i$ be a partition of the state set Q. Let $i := h+1$ with $h \in \mathbb{N}$ of Lem. 4.1.6. The following proves for all $q \in M_i$ that neither α_i nor β_i is able to reach a state outside M_i starting from q. Informally, M_i can be seen as the set of states which prevent a run of $\mathcal{D}_\varphi$, which takes too long to reach q_{acc}, from reaching q_{acc} at all.

Lemma 4.1.7. Let $w_{n,m}$ be a word of Def. 3.0.1, h of Lem. 4.1.6, $\gamma_i \in \{\alpha_i, \beta_i\}$ with $i := h+1$ and $(\text{reach}_{q_{\text{acc}}}(\alpha_i) \cup \text{reach}_{q_{\text{acc}}}(\beta_i)) \cup M_i$ a partition of Q. For all $q \in M_i$ and $q' \in Q$ with $q \in \text{reach}_{q'}(\gamma_i)$ it holds that $q' \in M_i$.

Proof. The proof is done for the case $\gamma_i = \alpha_i$. The case β_i is proven in the exact same way. Let $q \in M_i$ and $q \in \text{reach}_{q'}(\gamma_i)$. Most of the following arguments make use of the result of Lem. 4.1.6 which states that $\text{reach}_{q_{\text{acc}}}(\alpha_h) = \text{reach}_{q_{\text{acc}}}(\alpha_i)$ and $\text{reach}_{q_{\text{acc}}}(\beta_h) = \text{reach}_{q_{\text{acc}}}(\beta_i)$. This property is referred to as $(*)$.

Per definition it holds that $\alpha_i = \alpha_h^n \beta_h^m \alpha_h^n$. Let $p, p' \in Q$ be the states with $q \in \text{reach}_p(\alpha_h^{n-1})$ and $p \in \text{reach}_{p'}(\alpha_h)$. The first step is to argue that $p \in \text{reach}_{q_{\text{acc}}}(\beta_i) \setminus \text{reach}_{q_{\text{acc}}}(\alpha_i)$ or $p \in M_i$. This situation is depicted in Fig. 4.1 a).

It is shown that the case $p \in \text{reach}_{q_{\text{acc}}}(\alpha_i)$ leads to a contradiction. Assume that $p \in \text{reach}_{q_{\text{acc}}}(\alpha_i)$. With $(*)$ this implies that $q \in \text{reach}_{q_{\text{acc}}}(\alpha_i)$, which is a contradiction to the assumption that $q \in M_i$. Thus, it must hold that $p \in \text{reach}_{q_{\text{acc}}}(\beta_i) \setminus \text{reach}_{q_{\text{acc}}}(\alpha_i)$ or $p \in M_i$. The next step is to argue that this implies $p' \in M_i$ (see Fig. 4.1 b)).

Assume that p' is an element of $\text{reach}_{q_{\text{acc}}}(\alpha_i)$. It follows that $p \in \text{reach}_{q_{\text{acc}}}(\alpha_h^2)$, because of $p \in \text{reach}_{p'}(\alpha_h)$ and $(*)$. But this means that $p \in \text{reach}_{q_{\text{acc}}}(\alpha_i)$, because of $\alpha_i = \alpha_h^n \beta_h^m \alpha_h^n$. This is a contradiction to the findings about p. Thus, it must hold that $p' \in \text{reach}_{q_{\text{acc}}}(\beta_i) \setminus \text{reach}_{q_{\text{acc}}}(\alpha_i)$ or $p' \in M_i$. The case $p' \in \text{reach}_{q_{\text{acc}}}(\beta_i) \setminus \text{reach}_{q_{\text{acc}}}(\alpha_i)$ implies that $q \in \text{reach}_{q_{\text{acc}}}(\alpha_i)$. Again, this is a contradiction to $q \in M_i$. Thus, it follows that $p' \in M_i$. The next step is to argue that the state p'' with $p' \in \text{reach}_{p''}(\beta_h^m)$ is in M_i (see Fig. 4.1 c)).

The case $p'' \in \text{reach}_{q_{\text{acc}}}(\alpha_i)$ is a contradiction to $q \in M_i$, because of $(*)$. Thus, it holds that $p'' \in \text{reach}_{q_{\text{acc}}}(\beta_i) \setminus \text{reach}_{q_{\text{acc}}}(\alpha_i)$ or $p'' \in M_i$. Assume that $p'' \in \text{reach}_{q_{\text{acc}}}(\beta_i) \setminus \text{reach}_{q_{\text{acc}}}(\alpha_i)$. This implies that $p' \in \text{reach}_{q_{\text{acc}}}(\beta_h^m \beta_i)$. Then it also holds that $p' \in \text{reach}_{q_{\text{acc}}}(\beta_h^{m+1})$, because of $(*)$. But with $\beta_i = \beta_h^n \alpha_h^m \beta_h^n$ and $n > m$ follows that $p' \in \text{reach}_{q_{\text{acc}}}(\beta_i)$. This is a contradiction to $p' \in M_i$. Thus, it follows that $p'' \in M_i$. The last step in this proof is to argue that $q' \in M_i$ under the assumption that $p'' \in M_i$ (see Fig. 4.1 d)). It can be seen that the situation is essentially the same as in the step from state q to p'. This means that the arguments are the same and that it must hold that $q' \in M_i$. □

With these preliminary considerations, the key finding of this section becomes provable. Namely, on all suffixes of all $w_{n,m}$ the length of accepting runs of $\mathcal{D}_\varphi$ is bounded.

Theorem 4.1.8. Let $w_{n,m}$ be a word of Def. 3.0.1. There is $k \in \mathbb{N}$ such that for all $i_0 \in L_{w_{n,m}}(\mathcal{D}_\varphi)$ the length of the run π of $\mathcal{D}_\varphi$ on $w_{n,m}[i_0 :]$ is $|\pi| \leq k$.

Proof. Let $\pi = q_0\, q_1 \,\cdots\, q_{n'-1}$ be the accepting run of $\mathcal{D}_\varphi$ on $w_{n,m}[i_0 :]$ for some $i_0 \in L_{w_{n,m}}(\mathcal{D}_\varphi)$, $u = w[i_0 : (i_0 + (n'-1))]$ be the infix of $w_{n,m}$ that corresponds to π and $i := h + 1$ where h is the bound of Lem. 4.1.6. It is important to note that i is independent of π. For both cases $\alpha_i \notin \inf(u)$ and $\alpha_i \in \inf(u)$ it is shown that i induces a bound for $|u|$. This bound for the length of u implies a bound for the length of π.

If $\alpha_i \notin \inf(u)$ then it follows by the fourth observation of Prop. 3.0.3 that $|u| \leq (2+m)(2n+m)^i$.

If $\alpha_i \in \inf(u)$ then there are $v_{\text{pre}}, v_{\text{post}} \in \{a,b\}^*$ such that $u = v_{\text{pre}}\alpha_i v_{\text{post}}$ and $\alpha_i \notin v_{\text{pre}}$. From the previous case it is known that $|v_{\text{pre}}| \leq (2+m)(2n+m)^i$. Thus, it is left to argue that there is a bound for the length of v_{post}. From Prop. 3.0.4 it follows that there is $l \in \mathbb{N}$ such that $u \in \text{pre}(v_{\text{pre}}\, \alpha_i\, \gamma_0\, \gamma_1 \,\cdots\, \gamma_{l-1})$ with $\gamma_j \in \{\alpha_i, \beta_i\}$ for all $j < l$. The proof is proceeded by arguing that there is a bound, which is induced by i, for the minimum l needed to reach q_{acc}. This implies that $|v_{\text{post}}|$ is bounded as v_{post} is a prefix of $\gamma_0\, \gamma_1 \,\cdots\, \gamma_{l-1}$ for any l that reaches q_{acc}. Consider the partition $(\text{reach}_{q_{\text{acc}}}(\alpha_i) \cup \text{reach}_{q_{\text{acc}}}(\beta_i)) \cup M_i$ of Lem. 4.1.7. Let q_j be the state of π after v_{pre}. The proof is proceeded with a case distinction on q_j in context of this partition.

If $q_j \in M_i$ then it follows with Lem. 4.1.7 that neither with α_i nor with β_i the state q_j can reach q_{acc}. This means that there is no l such that the run corresponding to $v_{\text{pre}}\, \alpha_i\, \gamma_0\, \gamma_1 \,\cdots\, \gamma_{l-1}$ reaches q_{acc}. This is a contradiction, as π is accepting and u is a prefix of $v_{\text{pre}}\, \alpha_i\, \gamma_0\, \gamma_1 \,\cdots\, \gamma_{l-1}$ for some l.

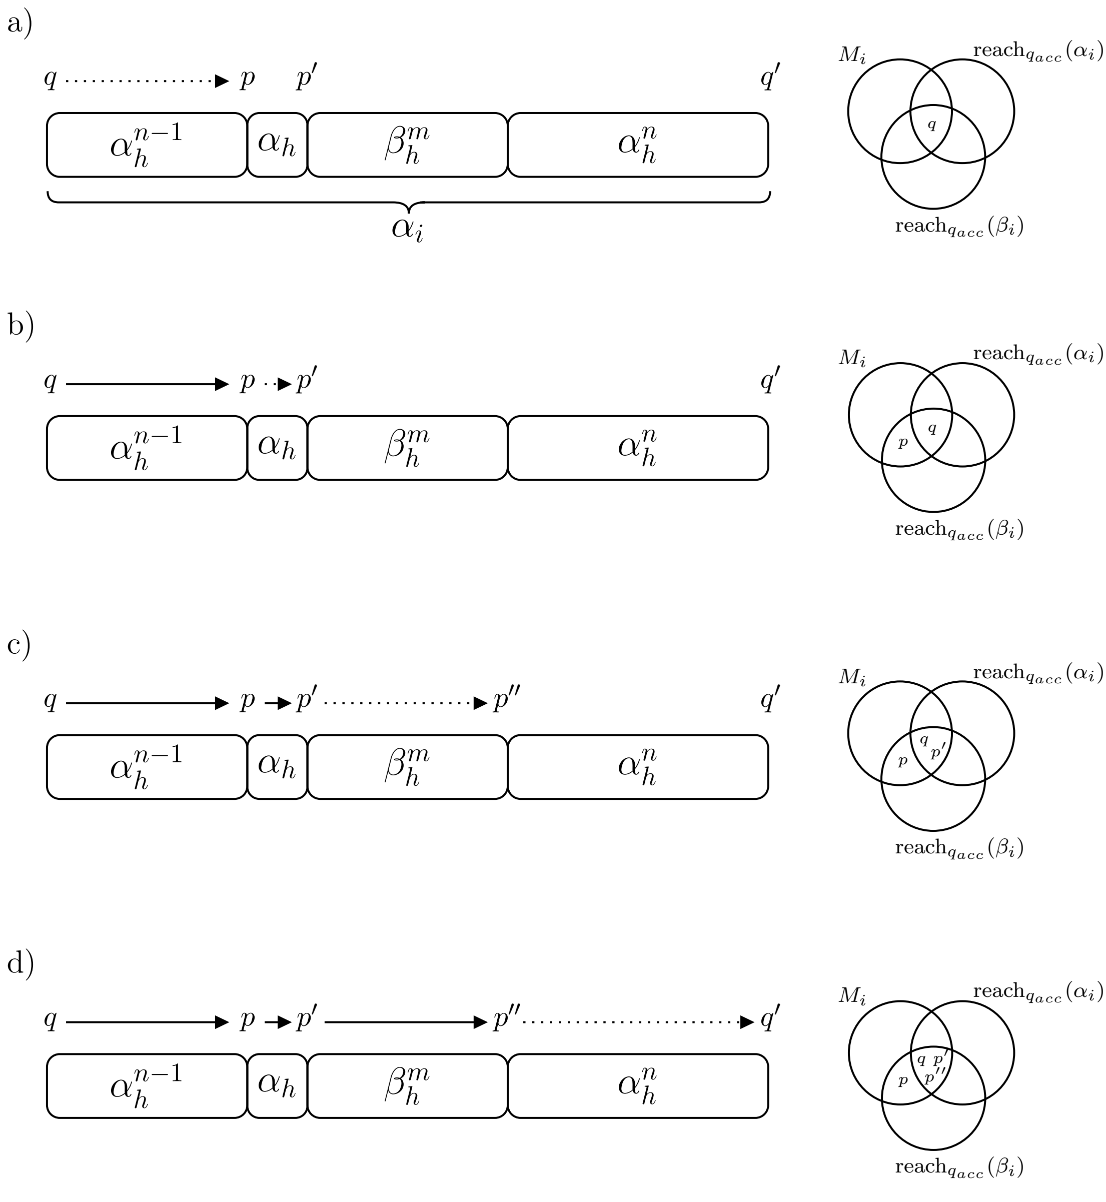

FIGURE 4.1: Visualisation of the proof steps of Lem. 4.1.7. The following things are depicted for each step: The left side shows the transition between states of $\mathcal{D}_\varphi$ currently under consideration (dotted arrow) in connection with the position in word α_i. The right side shows which assumptions are made about the considered states of $\mathcal{D}_\varphi$ in relation to the partition $(\text{reach}_{q_{\text{acc}}}(\alpha_i) \cup \text{reach}_{q_{\text{acc}}}(\beta_i)) \cup M_i$.

For the case that $q_j \in \text{reach}_{q_{\text{acc}}}(\alpha_i)$ it holds that $l = 0$.

If $q_j \in \text{reach}_{q_{\text{acc}}}(\beta_i) \setminus \text{reach}_{q_{\text{acc}}}(\alpha_i)$ then consider the state $q_{j'}$ after $v_{\text{pre}}\alpha_i$. That $q_{j'} \in \text{reach}_{q_{\text{acc}}}(\alpha_i)$ is not possible as it implies that $q_j \in \text{reach}_{q_{\text{acc}}}(\alpha_i^2)$ and thus $q_j \in \text{reach}_{q_{\text{acc}}}(\alpha_{i+1})$. As it holds that $q_j \notin \text{reach}_{q_{\text{acc}}}(\alpha_i)$, this is a contradiction to observation two of Lem. 4.1.6. The case $q_{j'} \in M_i$ is also not possible. The argument is the same as for the case $q_j \in M_i$. If $q_{j'} \in \text{reach}_{q_{\text{acc}}}(\beta_i) \setminus \text{reach}_{q_{\text{acc}}}(\alpha_i)$ consider γ_0. If $\gamma_0 = \beta_i$ it holds that $l = 1$. If $\gamma_0 = \alpha_i$ the argument is repeated for the state $q_{j''}$ after $v_{\text{pre}}\alpha_i\alpha_i$. With the help of observation five of Prop. 3.0.3 it follows that $l \leq 2n$. This bound respects the worst case, namely that all γ_k with $k < 2n$ are equal to α_i and $\gamma_{2n} = \beta_i$.

This shows for each case of q_j that v_{post} is bounded and therefore for both cases $\alpha_i \notin \inf(u)$ and $\alpha_i \in \inf(u)$ it proves that $|u|$ is bounded. This implies that $|\pi|$ is bounded and as i is independent of π it means that there is a bound k for all accepting runs of $\mathcal{D}_\varphi$ on all suffixes of $w_{n,m}$. □

So far, the combination of the results of this section shows for all $w_{n,m}$ that there is an unfolding φ^k which is semantically equivalent to φ. But it still needs to be shown that this implies the finite convergence of φ for each $w_{n,m}$ and that guardedness is no restriction to the validity of these findings for unguarded sentences.

From Boundedness of Automata to Finite Convergence of Unipolar Formulas

Theorem 4.1.9. Let $\varphi \in \mathcal{L}_\mu^{\text{lin}}$ be a guarded, μ-unipolar sentence and $w_{n,m}$ a word of Def. 3.0.1. It holds that $\varphi \in \text{FC}(w_{n,m})$.

Proof. The starting point for this proof is Th. 4.1.8. This theorem states that there is a bound m' for all $i \in L_{w_{n,m}}(\mathcal{D}_\varphi)$ on the length of the accepting run of $\mathcal{D}_\varphi$ on $w_{n,m}[i:]$. Then, Lem. 4.1.5 states that $L_{w_{n,m}}(\mathcal{D}_\varphi) = L_{w_{n,m}}(\mathcal{N}_\varphi)$ and that m' is a bound for all $i \in L_{w_{n,m}}(\mathcal{N}_\varphi)$ on at least one accepting run of $\mathcal{N}_\varphi$ on $w_{n,m}[i:]$. Then, Lem. 4.1.3 states that $L_{w_{n,m}}(\mathcal{N}_\varphi) = L_{w_{n,m}}(\mathcal{A}_\varphi)$ and that $m' - 1$ is a bound for all $i \in L_{w_{n,m}}(\mathcal{A}_\varphi)$ on the height of at least one accepting run of $\mathcal{A}_\varphi$ on $w_{n,m}[i:]$. With Lem. 4.1.1 this leads to the result that $[\![\varphi^{m'}]\!]^w_\eta = [\![\varphi]\!]^w_\eta$. Fix m' to be such unfolding for φ.

It needs to be argued that there is $n' \in \mathbb{N}$ for each $\mu x.\, \psi \in \text{sub}(\varphi)$ such that $[\![(\mu x.\, \psi)^{n'}]\!]^{w_{n,m}}_{\eta_x} = [\![\mu x.\, \psi]\!]^{w_{n,m}}_{\eta_x}$. For each subformula $\mu x_i.\, \psi_i$ it is shown that there exists n_i, satisfying the property $[\![(\mu x_i.\, \psi_i)^{n_i}]\!]^{w_{n,m}}_{\eta_{x_i}} = [\![\mu x_i.\, \psi_i]\!]^{w_{n,m}}_{\eta_{x_i}}$. Then, it follows with the monotonicity properties of Lem. 2.2.8 that the overall n is the maximum of all n_i.

For an outermost $\mu x_0.\, \psi_0$ fixpoint sentence it holds that $n_0 = m'$. Consider an inner fixpoint formula $\mu x_k.\psi_k$ and let $x_0, x_1, \ldots, x_{k-1}$ be a sequence of all mutually distinct fixpoint variables such that $\text{nd}(\mu x_i.\, \psi_i) \geq \text{nd}(\mu x_{i+1}.\psi_{i+1})$ for all $i < k$. It is assumed that $n_i \in \mathbb{N}$ is the unfolding such that $[\![(\mu x_i.\, \psi_i)^{n_i}]\!]^{w_{n,m}}_{\eta_{x_i}} = [\![\mu x_i.\, \psi_i]\!]^{w_{n,m}}_{\eta_{x_i}}$. Thus, it follows that

$$[\![\mu x_k.\psi_k[(\mu x_{k-1}.\psi_{k-1})^{n_{k-1}}/x_{k-1}] \cdots [(\mu x_0.\, \psi_0)^{n_0}/x_0]]\!]^{w_{n,m}}_\eta = [\![\mu x_k.\psi_k]\!]^{w_{n,m}}_{\eta_{x_k}}.$$

Then, with the result of Th. 4.1.8 and corresponding preliminary steps, there is an unfolding n_k such that $[\![(\mu x_k.\psi_k)^{n_k}]\!]^{w_{n,m}}_{\eta_{x_k}} = [\![\mu x_k.\psi_k]\!]^{w_{n,m}}_{\eta_{x_k}}$. □

Using this result for guarded, μ-unipolar sentences, it can be argued that the same holds for guarded, ν-unipolar sentences. Let φ be a ν-unipolar sentence. The idea is to argue that each rejecting run of $\mathcal{A}_\varphi$ on any suffix of some $w_{n,m}$ is bounded. To ensure the existence of rejecting runs, the automaton $\mathcal{A}_\varphi$ is extended by a rejecting state q_{rej} and then simplified as described at the end of Sec. 2.4.1. The translation into a nondeterministic and then into a deterministic automaton as well as the following reachability arguments are similar, respecting the state q_{rej} instead of q_{acc}. The monotonicity argument of Lem. 2.2.8 is the same just in the vice-versa direction.

Corollary 4.1.10. Let $\varphi \in \mathcal{L}_\mu^{\text{lin}}$ be a guarded, unipolar sentence and $w_{n,m}$ a word of Def. 3.0.1. It holds that $\varphi \in \text{FC}(w_{n,m})$.

Now, it remains to extend the result from a guarded, unipolar sentence to a possibly unguarded, unipolar sentence. This is done by applying the guarding procedure presented in Sec. 2.2.5 in a backwards direction, while arguing that finite convergence is preserved in each step.

Theorem 4.1.11. Let $\varphi \in \mathcal{L}_\mu^{\text{lin}}$ be an unipolar sentence and $w_{n,m}$ a word of Def. 3.0.1. It holds that $\varphi \in \text{FC}(w_{n,m})$.

Proof. Let φ be σ-unipolar. If φ is guarded the result is given by Cor. 4.1.10. If φ is unguarded then consider the guarded, σ-unipolar sentence φ' which is the result of the guarding procedure of Sec. 2.2.5 applied on φ. Again, with Cor. 4.1.10 it follows that φ' is finitely converging over $w_{n,m}$. Assume that all due to Prop. 2.2.13 generated $\hat{\sigma}$ in a subformula $\sigma x_i.\, \psi_i \in \text{sub}(\varphi')$ are identifiable and exclusively referred to as $\hat{\sigma_i}$. Furthermore, assume that all subformulas $\psi[\sigma x.\, \psi/x] \in \text{sub}(\varphi')$ generated using Prop. 2.2.12 are identifiable and referred to as unrolled subformulas. It is straightforward that applying the guarding procedure in the opposite direction on φ' results in φ. Thus, it needs to be argued that each reverse step of the procedure does preserve finite convergence. Let φ'_{tmp} denote an intermediate form in the reverse procedure.

First, consider replacing all $\hat{\sigma_i}$ with x_i in a subformula $\sigma x_i.\, \psi_i$ of the current sentence φ'_{tmp}. This corresponds to undoing the replacement of all not weakly-guarded occurrences of a fixpoint variable. Assume that m' is a point of finite convergence of φ'_{tmp} over $w_{n,m}$. For all $k \in \mathbb{N}$ can be inferred with a straightforward induction that $[\![(\sigma x_i.\, \psi_i)^{m'[k/c_i]}]\!]^w_\eta = [\![(\sigma x_i.\, \psi_i[x_i/\hat{\sigma_i}])^{m'[k/c_i]}]\!]^w_\eta$. This is a result from the fact that $\hat{\sigma_i}$ originally replaced not weakly-guarded occurrences of x_i. It follows from this that $[\![(\sigma x_i.\, \psi_i)^{m'}]\!]^w_\eta = [\![(\sigma x_i.\, \psi_i[x_i/\hat{\sigma_i}])^{m'}]\!]^w_\eta$. Furthermore, from Prop. 2.2.13 it follows for each $\sigma x_j.\, \psi_j \in \text{sub}(\psi_i)$ that it still holds $[\![(\sigma x_j.\, \psi_j)^{m'}]\!]^{w_{n,m}}_{\eta_{x_j}} = [\![\sigma x_j.\, \psi_j]\!]^{w_{n,m}}_{\eta_{x_j}}$ after replacing $\hat{\sigma_i}$ with x_i. In conclusion, undoing the replacement of not weakly-guarded x_i preserves the point of finite convergence m'.

Consider now an unrolled formula $\psi_i[\sigma x_i.\, \psi_i/x_i]$ and assume that m' is a point of finite convergence of the current sentence φ'_{tmp}. It needs to be argued that $\psi_i[\sigma x_i.\, \psi_i/x_i]$ can be replaced with $\sigma x_i.\, \psi_i$ such that finite convergence is preserved. This corresponds to undoing the unfolding of inner fixpoint subformulas. Note, φ'_{tmp} is a σ-unipolar sentence. With the monotonicity properties of Lem. 2.2.8 it follows that $m'+1$ also is a point of finite convergence of φ'_{tmp}. Let $\eta^{m'}_{x_i}$ be the environment that maps each fixpoint variable $x_h \in \text{sub}(\psi_i)$ occurring freely in ψ_i to

$$[\![(\sigma x_h.\psi_h)^{(m'+1)[m'/c_h]}]\!]^{w_{n,m}}_{\eta^{m'}_{x_h}}.$$

The environments $\eta^{m'}_{x_i}$ and $\eta^{m'}_{x_h}$ are semantic equivalents to the syntactic replacement of free occurring fixpoint variables with the m' unfolding of their respective fixpoint subformula. From the monotonicity properties of Lem. 2.2.8 and the fact that m' is a point of finite convergence of φ'_{tmp} over $w_{n,m}$ it follows that

$$[\![(\sigma x_h.\psi_h)^{(m'+1)[m'/c_h]}]\!]^{w_{n,m}}_{\eta^{m'}_{x_h}} = [\![\sigma x_h.\psi_h]\!]^{w_{n,m}}_{\eta_{x_h}}.$$

As $m'+1$ also is a point of finite convergence of φ'_{tmp} over $w_{n,m}$, it holds for all $\sigma x_j.\,\psi_j \in \text{sub}(\psi_i)$ that

$$[\![(\sigma x_j.\,\psi_j)^{m'+1}]\!]^{w_{n,m}}_{\eta_{x_j}} = [\![\sigma x_j.\,\psi_j]\!]^{w_{n,m}}_{\eta_{x_j}} \text{ which leads to}$$
$$[\![(\sigma x_j.\,\psi_j)^{m'+1}]\!]^{w_{n,m}}_{\eta^{m'}_{x_j}} = [\![\sigma x_j.\,\psi_j]\!]^{w_{n,m}}_{\eta^{m'}_{x_j}}.$$

The second equivalence states for each fixpoint subformula of ψ_i that the semantics of the $(m'+1)$-th unfolding are equal to the true semantics. This implies that

$$[\![(\psi_i[\sigma x_i.\,\psi_i/x_i])^{m'+1}]\!]^{w_{n,m}}_{\eta^{m'}_{x_i}} = [\![\psi_i[\sigma x_i.\,\psi_i/x_i]]\!]^{w_{n,m}}_{\eta^{m'}_{x_i}} \text{ and with Prop. 2.2.12 that}$$
$$[\![(\psi_i[\sigma x_i.\,\psi_i/x_i])^{m'+1}]\!]^{w_{n,m}}_{\eta^{m'}_{x_i}} = [\![(\sigma x_i.\,\psi_i)^{(m'+1)}]\!]^{w_{n,m}}_{\eta^{m'}_{x_i}}.$$

Thus, the unrolled formula $\psi_i[\sigma x_i.\,\psi_i/x_i]$ can be replaced with $\sigma x_i.\,\psi_i$ without changing the semantics of $\varphi'^{(m'+1)}_{\text{tmp}}$. That $m'+1$ is a point of finite convergence after this replacement follows from the fact that $\sigma x_i.\,\psi_i$ was a subformula of φ'_{tmp} before the replacement. Therefore, the subformula conditions of a finite convergence are still valid.

These two arguments show that each reverse step of the guarding procedure preserves finite convergence over $w_{n,m}$ which implies that φ is finitely converging over $w_{n,m}$. □

4.2 Finite Convergence of Alternation-Free Formulas

With the results of Sec. 4.1 about the finite convergence of unipolar sentences over all $w_{n,m}$, a proof of finite convergence of all alternation-free sentences over all $w_{n,m}$ becomes arguable. Informally put, the idea behind this proof is that alternation-free sentences are in fact a stack of unipolar sentences.

Theorem 4.2.1. Let $\varphi \in \mathcal{L}^{\text{lin}}_{\mu}$ be an alternation-free sentence and $w_{n,m}$ a word of Def. 3.0.1. It holds that $\varphi \in \text{FC}(w_{n,m})$.

Proof. Let V^{μ}_{φ} and V^{ν}_{φ} be the sets of least and greatest fixpoint variables occurring in φ. It follows from the fact that φ is alternation-free that there is no pair $x \in V^{\mu}_{\varphi}$ and $y \in V^{\nu}_{\varphi}$ such that x occurs free in $\text{sub}(\text{fp}_{\varphi}(y))$ or y occurs free in $\text{sub}(\text{fp}_{\varphi}(x))$. This property leads to the following ordering: Each unipolar fixpoint sentence $\sigma x.\,\psi \in \text{sub}(\varphi)$ is of order 0, denoted by $\text{order}(\sigma x.\,\psi) = 0$. For a non-unipolar sentence $\sigma x.\,\psi \in \text{sub}(\varphi)$ the order is defined as

$$\max(\{\text{order}(\sigma' x'.\,\psi') \mid \sigma' x'.\,\psi' \in \text{sub}(\psi) \text{ is a sentence and } \sigma \neq \sigma'\}) + 1.$$

If $\text{order}(\sigma x.\,\psi) = 0$ the sentence is unipolar. For this case it follows by Th. 4.1.11 that $\sigma x.\,\psi$ is finitely converging over $w_{n,m}$.

For the induction step let $\sigma x.\, \psi \in \mathrm{sub}(\varphi)$ be a sentence and let there be k sentences $\sigma_i x_i.\psi_i \in \mathrm{sub}(\psi)$ with $\mathrm{order}(\sigma_i x_i.\psi_i) = \mathrm{order}(\sigma x.\, \psi) - 1$. The induction hypothesis states that each of these subsentences is finitely converging over $w_{n,m}$. Let n_i be the corresponding unfolding. Then it follows that

$$[\![\sigma x.\, \psi]\!]_\eta^{w_{n,m}} = [\![\sigma x.\, \psi[(\sigma_i x_i.\, \psi_i)^{n_i}/\sigma_i x_i.\, \psi_i]_0^{k-1}]\!]_\eta^{w_{n,m}},$$

where $[\cdots]_0^{k-1}$ denotes the replacement for all $0 \leq i \leq k$. This new sentence is unipolar. Thus, with Th. 4.1.11 it follows that there is some n' which is a point of finite convergence for $\sigma x.\, \psi[(\sigma_i x_i.\, \psi_i)^{n_i}/\sigma_i x_i.\, \psi_i]_0^{k-1}$ over $w_{n,m}$. As each replaced fixpoint subformula is an unipolar sentence, the monotonicity properties of Lem. 2.2.8 can be applied. Then, the point of finite convergence of $\sigma x.\, \psi$ over $w_{n,m}$ is given by $\max(\{n_0, \ldots, n_{k-1}, n'\})$.

Using this, it can be assumed that each outermost fixpoint sentence of φ is finitely converging over $w_{n,m}$. An additional induction argument over the structure of φ gives the final result. □

4.3 Presumed Generalisation of the Result

This section sketches a conjecture on how to use the proven results of this work to show that all $\mathcal{L}_\mu^{\mathrm{lin}}$-formulas are finitely converging over all $w_{n,m}$. The main result to fall back on is given by Th. 4.2.1. Namely, the result that all alternation-free $\mathcal{L}_\mu^{\mathrm{lin}}$-formulas are finitely converging over all $w_{n,m}$.

The idea of the proof is informally described as follows: Consider a formula $\varphi \in \mathcal{L}_\mu^{\mathrm{lin}}$ and a word $w_{n,m}$. It is shown for a fixpoint subformula $\sigma x.\psi$ that there is a semantically equivalent finite unfolding while assuming that outer fixpoint subformulas are unfolded already and that inner fixpoint subformulas are computed completely. Applying this procedure in a top-down manner to each fixpoint subformula should lead to a proof for the finite convergence of φ over $w_{n,m}$. How to argue that there is a semantically equivalent unfolding for an outermost fixpoint formula, while computing inner fixpoints completely, is shown in the following. This is followed by a brief discussion of the difficulties that come with this approach. These difficulties are one of the reasons why a formal proof is not given in this work.

First, some preliminaries from the theory of basic modal logic are needed. Basic modal logic is not formally defined in Chap. 2, but can be regarded as the fragment of $\mathcal{L}_\mu$-formulas φ such that $\mathrm{ad}(\varphi) = 0$ and that there are no free occurring variables in φ. This also implies that there is no subformula $\sigma x.\psi \in \mathrm{sub}(\varphi)$. The same holds for the linear-time variant $\mathcal{L}_\mu^{\mathrm{lin}}$. For the remaining part of this section a formula $\varphi \in \mathcal{L}_\mu^{\mathrm{lin}}$ is said to be a modal logic formula if $\mathrm{ad}(\varphi) = 0$. For further insights into the theory of modal logic it is referred to [BBW07]. The first observation needed is that the inclusion of a position i of an infinite word $w \in \Sigma^\omega$ in the semantics $[\![\varphi]\!]_\eta^w$ of a basic modal logic formula with $\mathrm{md}(\varphi) = d$ is only dependent on the first d positions after i. This leads to the case that either all i for which $w[i : i + d]$ looks the same are included in the semantics of $[\![\varphi]\!]_\eta^w$ or none of them.

Observation 4.3.1. Let φ be modal logic formula with $\mathrm{md}(\varphi) = d$, $w \in \Sigma^\omega$ and $I \subseteq \mathbb{N}$ such that for all $i, i' \in I$ it holds $w[i : i + d] = w[i' : i' + d]$. It holds that $I \subseteq [\![\varphi]\!]_\eta^w$ or $I \cap [\![\varphi]\!]_\eta^w = \emptyset$.

The second observation needed is an implication of this understanding. Across all $w \in \Sigma^\omega$ each modal logics formula φ is semantically equivalent to a disjunction $\bigvee_{u \in U} \varphi_u$. Here, U is the set of all finite words of length d that satisfy φ and the formula φ_u is satisfied by a position i if and only if the infix of length d starting in i is equal to u. Additionally, it can be shown that U is finite.

Observation 4.3.2. Let φ be a modal logics formula and $\bigvee_{u \in U} \varphi_u$ as described above. For all $w \in \Sigma^\omega$ it holds that $[\![\varphi]\!]^w_\eta = [\![\bigvee_{u \in U} \varphi_u]\!]^w_\eta$.

Consider now an (arbitrary) formula $\varphi \in \mathcal{L}_\mu^{\text{lin}}$ and w.l.o.g. assume that $\varphi = \mu x.\psi$. The following arguments are dual for the ν-case. Consider now the single unfolding of $\mu x.\psi$, defined as $(\mu x.\psi, 0) = \bot$ and $(\mu x.\psi, i+1) = \psi[(\mu x.\psi, i)/x]$. This is a syntactic variant of Def. 2.2.5 for finite steps. For the remaining part of this section unfolding refers to this notion of a single fixpoint unfolding. It is straightforward that for all $w \in \Sigma^\omega$ it holds

$$[\![\varphi, 0]\!]^w_\eta \subseteq [\![\varphi, 1]\!]^w_\eta \subseteq \ldots \subseteq [\![\varphi]\!]^w_\eta.$$

It is a well-known result by Kaivola that for each $\mathcal{L}_\mu^{\text{lin}}$-formula there is a semantically equivalent alternation-free formula over all infinite words [Kai95]. With the results of this work, particularly Th. 4.2.1, it follows that there is also an equivalent modal logic formula over each $w_{n,m}$. Fix some word $w_{n,m}$. Let φ' be the equivalent modal logic formula for φ over $w_{n,m}$ and assume that $\text{md}(\varphi') = d$. The arguments of Obs. 4.3.2 lead to the case that there is a set U of words of length d such that $[\![\varphi']\!]^{w_{n,m}}_\eta = [\![\bigvee_{u \in U} \varphi'_u]\!]^{w_{n,m}}_\eta$. With the same arguments it follows for each unfolding (φ, i) that there is a set U_i and with the relation $[\![\varphi, i]\!]^{w_{n,m}}_\eta \subseteq [\![\varphi]\!]^{w_{n,m}}_\eta$ it follows that $U_i \subseteq U$. This leads to the sequence of inclusions $U_0 \subseteq U_1 \subseteq \ldots \subseteq U$. It is known that U is finite. This means that this sequence stabilizes with some U_i, which implies that $[\![(\varphi, i)]\!]^w_\eta = [\![\varphi]\!]^w_\eta$.

Observation 4.3.3. Let $\varphi \in \mathcal{L}_\mu^{\text{lin}}$ and let $w_{n,m}$ a word of Def. 3.0.1. There is $i \in \mathbb{N}$ such that $[\![(\varphi, i)]\!]^{w_{n,m}}_\eta = [\![\varphi]\!]^{w_{n,m}}_\eta$.

At this point it can be inferred that for each $w_{n,m}$ and $\mathcal{L}_\mu^{\text{lin}}$-formula φ there is a semantically equivalent unfolding of all outermost fixpoint formulas (φ, i). It is left to argue that this implies the finite convergence of φ over all $w_{n,m}$. This step is the one of the reasons why a general proof is only sketched in this work: The validity of the previous arguments is not doubted, but using these in an induction over all fixpoint subformulas of φ comes with difficulties.

Consider an inner fixpoint subformula $\sigma x.\psi$ of φ and assume that there are semantically equivalent unfoldings for outer fixpoint subformulas. The same procedure as described above can be applied to $\sigma x.\psi$ while assuming that all fixpoint variables x' which are occurring freely in $\text{sub}(\psi)$ are replaced with the respective unfolding of their defining fixpoint subformula $\sigma' x'.\psi'$. However, the unfolding of $\sigma' x'.\psi'$ includes fixpoint subformulas, especially the subformula $\sigma x.\psi$, which means that $\sigma x.\psi$ is duplicated while replacing x' in ψ. As the unfolding of $\sigma x.\psi$ is done in reference to these outer unfoldings further arguments are needed to find one unfolding of $\sigma x.\psi$ for all occurrences. Because of this difficulty and the reliance on basic modal logic, which is not topic of this work, a formal proof is out of scope.

Conjecture 4.3.4. Let $\varphi \in \mathcal{L}_\mu^{\text{lin}}$ be a sentence and $w_{n,m}$ a word of Def. 3.0.1. It holds that $\varphi \in \text{FC}(w_{n,m})$.

Chapter 5

HFL-Formulas without Finite Convergence

This chapters shows that there are $\mathrm{HFL}_{\mathrm{lin}}$-sentences, which are not finitely converging over some $w_{n,m}$. An example from the second-order fragment $\mathrm{HFL}^2_{\mathrm{lin}}$ and an example from the first-order fragment $\mathrm{HFL}^1_{\mathrm{lin}}$ is presented. Both example sentences are considered in context of the infinite word $w_{2,1}$. This means that all α_i, β_i in this chapter refer to the respective infix of $w_{2,1}$.

The $\mathrm{HFL}^2_{\mathrm{lin}}$-sentence φ_2 is considered first as its non-finite convergence is easier to see. The definition of φ_2 uses the abbreviations presented in Sec. 2.3.2:

$$\begin{aligned}
\varphi_2 = &\big(\nu x.\, \lambda f_\alpha, f_\beta.\, f_{\alpha+1} \top \wedge x\, f_{\alpha+1}\, f_{\beta+1}\big)\, (a \rightarrow)\, (b \rightarrow),\\
&(\psi \rightarrow) := \lambda y.\, \psi \wedge \odot\, y,\\
&f_{\alpha+1} := f_\alpha \circ f_\alpha \circ f_\beta \circ f_\alpha \circ f_\alpha,\\
&f_{\beta+1} := f_\beta \circ f_\beta \circ f_\alpha \circ f_\beta \circ f_\beta.
\end{aligned}$$

Consider φ_2 in context of $w_{2,1}$. Informally, a position i of $w_{2,1}$ is included in $[\![\varphi_2]\!]^{w_{2,1}}_\eta$ if all α_i start in this position. To understand this, finite unfoldings of φ_2 are considered. These are depicted in a simplified version, using β-reduction and the semantic equivalence of $\odot\,(\psi_1 \wedge \psi_2)$ and $\odot\,\psi_1 \wedge \odot\,\psi_2$. A subformula $(x_0 \rightarrow (x_1 \rightarrow \cdots (x_{n-1} \rightarrow \psi) \cdots)$ is denoted by $u \rightarrow \psi$ with $u = x_0 \cdots x_{n-1}$. If $\psi = \top$ it is further simplified and denoted by u. The first three unfoldings and the i-th one are given by

$$\begin{aligned}
\varphi_2^0 &= \top,\\
\varphi_2^1 &= \underbrace{aabaa}_{\alpha_1} \wedge \top,\\
\varphi_2^2 &= \underbrace{aabaaaabaabbabbaabaaaabaa}_{\alpha_2} \wedge \underbrace{aabaa}_{\alpha_1} \wedge \top \text{ and}\\
\varphi_2^i &= \alpha_i \wedge \cdots \wedge \alpha_1 \wedge \top.
\end{aligned}$$

Taking a look at Def. 3.0.1 it follows that no position of $w_{2,1}$ satisfies this property, which means that $[\![\varphi]\!]^{w_{2,1}}_\eta = \emptyset$ for all η. In contrast to the results about $\mathcal{L}^{\mathrm{lin}}_\mu$-sentences of Chap. 4 the $\mathrm{HFL}^2_{\mathrm{lin}}$-sentence φ_2 is not finitely converging over $w_{2,1}$. The reason for this is the stepwise inclusion of greater α_i, if seen as simplifications of $\rightarrow$ applications, in greater unfoldings of φ_2. This excludes further positions in each further unfolding.

Theorem 5.0.1. Let φ_2 be defined as above. It holds that $\varphi_2 \notin \mathrm{FC}(w_{2,1})$.

Proof. For all i it is shown that $[\![\varphi_2^i]\!]_\eta^{w_{2,1}} \neq [\![\varphi_2^{i+1}]\!]_\eta^{w_{2,1}}$. As there is only one fixpoint subformula present in φ_2, this implies that φ_2 is not finitely converging over $w_{2,1}$. Let $n_i = (1 - 5^i)/(1 - 5)$, which is the starting point of the first α_i in $w_{2,1}$. The following shows that for all i it holds that $n_i \in [\![\varphi_2^i]\!]_\eta^{w_{2,1}}$ and that there is no $j < n_i$ such that $j \in [\![\varphi_2^i]\!]_\eta^{w_{2,1}}$. This implies that $n_i \notin [\![\varphi_2^{i+1}]\!]_\eta^{w_{2,1}}$, which gives the desired statement. From the discussion above about the i-th unfolding of φ_2 it follows that $k \in [\![\varphi^i]\!]_\eta^{w_{2,1}}$ if each $\alpha_{i'}$ with $i' \leq i$ starts at position k of $w_{2,1}$. The second observation of Prop. 3.0.3 states that this is the case for position n_i as there starts α_i. Furthermore, this observation also implies that no position j with $j < n_i$ is included in $[\![\varphi^i]\!]_\eta^{w_{2,1}}$ as there starts no α_i-infix. This means that for all i it holds that $[\![\varphi_2^i]\!]_\eta^{w_{2,1}} \neq [\![\varphi_2^{i+1}]\!]_\eta^{w_{2,1}}$, which implies that $\varphi_2 \notin \mathrm{FC}(w_{2,1})$. □

This theorem shows that there are $\mathrm{HFL}^2_{\mathrm{lin}}$-sentences which are not finitely converging over some $w_{n,m}$. Besides this result, the sentence φ_2 is not that interesting as its semantics are trivial in context of $w_{2,1}$. This leaves room for critique that each meaningful $\mathrm{HFL}_{\mathrm{lin}}$-formula could be finitely converging over all $w_{n,m}$. Additionally, it could still be the case that all $\mathrm{HFL}^1_{\mathrm{lin}}$-formulas are finitely converging over all $w_{n,m}$. Nevertheless, φ_2 is a fairly simple example compared to the following $\mathrm{HFL}^1_{\mathrm{lin}}$-sentence.

The $\mathrm{HFL}^1_{\mathrm{lin}}$-formula φ_1 is considered in the followin. The definition of φ_1 makes use of the same abbreviations as the definition of φ_2:

$$\begin{aligned}
\varphi_1 &:= \nu x.\, \psi_\alpha\, x,\\
\psi_\alpha &:= \mu x_\alpha.\, \lambda y.\, (a \to y) \vee (x_\alpha\, (x_\alpha\, (\psi_\beta\, (x_\alpha\, (x_\alpha\, y))))),\\
\psi_\beta &:= \mu x_\beta.\, \lambda z.\, (b \to z) \vee (x_\beta\, (x_\beta\, (x_\alpha\, (x_\beta\, (x_\beta\, z))))).
\end{aligned}$$

Note, ψ_α and ψ_β are of type $\bullet \to \bullet$. The formula φ_1 is also considered in context of $w_{2,1}$. Among others reasons, a position of $w_{2,1}$ is included in the semantics of φ_1 if an infinite chain of arbitrary α_i starts from it. To show this, it can be inferred that for all i and $i' < i$ the unfolding ψ_α^i includes $\alpha_{i'} \to y$ as a disjunct, if $\alpha_{i'}$ is seen as a simplification of $\to$ applications. This can be proven with a straightforward induction. An intuitive understanding of the unfolding of φ_1 is given in the following. First, the first three and the i-th unfolding of ψ_α are considered. Disjuncts of an unfolding ψ_α^i, which are not of interest for this understanding, are subsumed in Ψ_α^i:

$$\psi_\alpha^0 = \lambda y.\bot,$$

$$\begin{aligned}
\psi_\alpha^1 &= \lambda y.(a \to y) \vee (\psi_\alpha^0\, (\psi_\alpha^0\, ((\psi_\beta[\psi_\alpha^0/x_\alpha])^1\, (\psi_\alpha^0\, (\psi_\alpha^0\, y)))))\\
&= \lambda y.(a \to y) \vee \bot,\\
\big(\psi_\beta[\psi_\alpha^0/x_\alpha]\big)^1 &= \lambda z.(b \to z) \vee \bot,
\end{aligned}$$

$$\begin{aligned}
\psi_\alpha^2 &= \lambda y.(a \to y) \vee (\psi_\alpha^1\, (\psi_\alpha^1\, ((\psi_\beta[\psi_\alpha^1/x_\alpha])^2\, (\psi_\alpha^1\, (\psi_\alpha^1\, y)))))\\
&= \lambda y.(a \to y) \vee (aabaa \to y) \vee \Psi_\alpha^2,\\
\big(\psi_\beta[\psi_\alpha^1/x_\alpha]\big)^2 &= \lambda z.(b \to z) \vee (bbabb \to z) \vee \bot,
\end{aligned}$$

$$\begin{aligned}
\psi_\alpha^i &= \lambda y.(a \to y) \vee (\psi_\alpha^{i-1}\, (\psi_\alpha^{i-1}\, ((\psi_\beta[\psi_\alpha^{i-1}/x_\alpha])^i\, (\psi_\alpha^{i-1}\, (\psi_\alpha^{i-1}\, y)))))\\
&= \lambda y.(\alpha_0 \to y) \vee \cdots \vee (\alpha_{i-1} \to y) \vee \Psi_\alpha^i.
\end{aligned}$$

As stated above, the i-th unfolding of ψ^i_α includes a disjunct $\alpha_{i'} \to y$ for all $i < i$. Next, the first three and the i-th unfolding of φ_1 are considered:

$$\begin{aligned}
\varphi_1^0 &= \top, \\
\varphi_1^1 &= \psi^1_\alpha\, \top, \\
\varphi_1^2 &= \psi^2_\alpha\, (\psi^2_\alpha\, \top) \text{ and} \\
\varphi_1^i &= \underbrace{\psi^i_\alpha\, (\psi^i_\alpha\, (\cdots (\psi^i_\alpha}_{i \text{ times}}(\top))\cdots)).
\end{aligned}$$

This shows that with each unfolding of φ_1 there occurs one more unfolding of ψ_α. In conclusion, the intuitive understanding gained so far can be described as follows. The i times unfolding of φ_1 creates a chain of length i of $\bullet \to \bullet$ building blocks taken from ψ^i_α which are ultimately applied to $\top$. The i times unfolding of ψ_α provides building blocks of the form $(\alpha_{i'} \to)$ for all $i' < i$. In combination, a position of $w_{2,1}$ is in the semantics of φ^i_1 if there starts a sequence of length i of $\alpha_{i'}$ such that $i' < i$. Therefore, if i goes to infinity, which corresponds to the full computation of the semantics of φ_1, a position of $w_{2,1}$ is in $[\![\varphi_1]\!]^{w_{2,1}}_\eta$ if an infinite chain of arbitrary α_i starts from it.

This understanding leads to the first result about φ_1 and $w_{2,1}$. In contrast to φ_2, the semantics $[\![\varphi_1]\!]^{w_{2,1}}_\eta$ define a non-trivial set. In particular, $[\![\varphi_1]\!]^{w_{2,1}}_\eta$ includes the starting points of all first occurring α_i. The reason for this is that at each of these points starts an infinite chain of α_i.

Theorem 5.0.2. Let φ_1 be defined as above and $w_{2,1}$ of Def. 3.0.1. For all $n_i = (1-5^i)/(1-5)$ it holds that $n_i \in [\![\varphi_1]\!]^{w_{2,1}}_\eta$. Furthermore, it holds that $[\![\varphi_1]\!]^{w_{2,1}}_\eta \neq \mathbb{N}$.

Proof. Let $x_0\, x_1\, \cdots$ with $x_i \in \{a, b\}$ be equal to $w_{2,1}$. For all $n_i \in \mathbb{N}$ it it is proven that $n_i \in [\![\varphi_1]\!]^{w_{2,1}}_\eta$. Informally put, it is shown that the starting point of each first α_i in $w_{2,1}$ is included in the semantics of φ_1 over $w_{2,1}$. Therefore, the formula ψ_α is considered. Let $I \subseteq \mathbb{N}$ be some subset of positions in $w_{2,1}$. From the understanding gained of finite unfoldings of ψ_α it follows for all $i \in I$ and all $j \in \mathbb{N}$ that if there is an infix $x_j\, \alpha_k[1:]\, x_i \in \inf(w_{2,1})$ for some α_k then $j \in [\![\psi_\alpha\, x]\!]^{w_{2,1}}_{\eta[x \mapsto I]}$. With this observation it can be argued that each n_i is included in $[\![\varphi_1]\!]^{w_{2,1}}_\eta$. By induction it is shown for each iteration step i of the greatest fixpoint (see Def. 2.2.5) that $n_i \in [\![\varphi_1, i]\!]^{w_{2,1}}_\eta$. It is a straightforward argument that $n_i \in [\![\varphi_1, 0]\!]^{w_{2,1}}_\eta$ is valid for all i. Consider $[\![\varphi_1, j+1]\!]^{w_{2,1}}_\eta$ and assume that $n_i \in [\![\varphi_1, j]\!]^{w_{2,1}}_\eta$ for all i. Fix some n_i and note that

$$[\![\varphi_1, j+1]\!]^{w_{2,1}}_\eta = [\![\psi_\alpha\, x]\!]^{w_{2,1}}_{\eta[x \mapsto [\![\varphi_1, j]\!]^{w_{2,1}}_\eta]}.$$

Per definition of n_i it follows that $x_{n_i}\, \alpha_i[1:]\, x_{n_{i+1}} \in \inf(w_{2,1})$. From the induction hypothesis it can be inferred that $n_{i+1} \in [\![\varphi_1, j]\!]^{w_{2,1}}_\eta$. With the understanding gained of unfoldings of ψ_α it follows that $n_i \in [\![\psi_\alpha\, x]\!]^{w_{2,1}}_{\eta[x \mapsto [\![\varphi_1, j]\!]^{w_{2,1}}_\eta]} = [\![(\varphi_1, j+1)]\!]^{w_{2,1}}_\eta$. Thus, it follows that each n_i is in each finite approximation of φ_1 and therefore in $\bigcap_{i \in \mathbb{N}} [\![(\varphi_1, i)]\!]^{w_{2,1}}_\eta$. With reference to the discussion in Sec. 2.5 it can be inferred that this intersection is equal to $[\![\varphi_1]\!]^{w_{2,1}}_\eta$ as $w_{2,1}$ is an infinite word. This gives the desired result.

That $[\![\varphi_1]\!]^{w_{2,1}}_\eta \neq \mathbb{N}$ holds can be seen in Conj. 5.0.4. It is stated there that for all i it follows that m_{i+1}, which is the starting point of the second α_i in the first α_{i+1} of $w_{2,1}$, is not included in $[\![\varphi_1]\!]^{w_{2,1}}_\eta$. □

Note, the positions n_i are not the only positions of $w_{2,1}$ included in $[\![\varphi_1]\!]_\eta^{w_{2,1}}$ as there are disjuncts included in the unfolding of φ_1 which do not follow the pattern of some α_j.

Another important implication of the understanding gained of the unfolding of φ_1 and ψ_α is the following: The position m_{i+1}, which is the starting point of the second α_i in the first α_{i+1}, is included in the semantics of the φ_1^i. The reason for this is that in m_{i+1} starts a sequence $\alpha_{i-1}\,\alpha_{i-2}\,\cdots\,\alpha_0$. This sequence also occurs as a disjunct in φ_1^i, if considered as a simplification of $\rightarrow$ applications.

Observation 5.0.3. Let φ_1 be defined as above and $w_{2,1}$ of Def. 3.0.1. For all $m_{i+1} = (1-5^{i+1})/(1-5) + (2n+m)^i$ it holds that $m_i \in [\![\varphi_1^i]\!]_\eta^{w_{2,1}}$.

In contrast to this, it can be shown that m_{i+1} is not included in the semantics of φ_1^{i+1} or any greater unfolding of φ_1. As a formal proof tends to be lengthy and tedious, it is not given in this work.

Conjecture 5.0.4. Let φ_1 be defined as above and $w_{2,1}$ of Def. 3.0.1. For all i and $j > i$ it holds that $m_{i+1} = (1-5^{i+1})/(1-5) + (2n+m)^i$ is not included in $[\![\varphi_1^j]\!]_\eta^{w_{2,1}}$.

From the combination of Obs. 5.0.3 and Con. 5.0.4 it can be inferred for all i that there is a position included in the semantics of φ_1^i but is not included in the semantics of φ. Thus, no finite unfolding φ_1^i is semantically equal to φ_1 over $w_{2,1}$. This implies that φ_1 is not finitely converging over $w_{2,1}$.

Conjecture 5.0.5. Let φ_1 be defined as above and $w_{2,1}$ of Def. 3.0.1. It holds that $\varphi_1 \notin w_{2,1} \in \mathrm{FC}()$.

From Th. 5.0.2 and Conj. 5.0.4 it follows for all i that the position n_{i+1}, which is the starting point of the highest-level α_{i+1} in $w_{2,1}$, is included in $[\![\varphi_1]\!]_\eta^{w_{2,1}}$ and that the position m_{i+1}, which is the starting point of the second α_i in this α_{i+1}, is not included in $[\![\varphi_1]\!]_\eta^{w_{2,1}}$. In Sec. 4.3 it is discussed that over each $w_{n,m}$ each $\mathcal{L}_\mu^{\mathrm{lin}}$-formula is equal to some formula of basic modal logic. Such a modal logic formula has some modal-depth d which means that it is not able to distinguish between n_{i+1} and m_{i+1} if $|\alpha_i| \geq d$. This implies for all $\mathcal{L}_\mu^{\mathrm{lin}}$-formulas φ that there is an i such that $n_{i+1} \in [\![\varphi]\!]_\eta^{w_{2,1}}$ if and only if $m_{i+1} \in [\![\varphi]\!]_\eta^{w_{2,1}}$. As the results for φ_1 regarding the positions n_{i+1} and m_{i+1} are valid for all i, it follows that $[\![\varphi_1]\!]_\eta^{w_{2,1}}$ is not definable by a $\mathcal{L}_\mu^{\mathrm{lin}}$-formula.

Observation 5.0.6. Let $w_{2,1}$ be of Def. 3.0.1. There is no $\varphi \in \mathcal{L}_\mu^{\mathrm{lin}}$ such that $[\![\varphi]\!]_\eta^{w_{2,1}} = [\![\varphi_1]\!]_\eta^{w_{2,1}}$.

These examples, especially the sentence φ_1, show that there are non-trivial $\mathrm{HFL}_{\mathrm{lin}}$-formulas which are not finitely converging over some $w_{2,1}$. Furthermore, it was shown that the semantics of φ_1 over $w_{2,1}$ are not $\mathcal{L}_\mu^{\mathrm{lin}}$-definable. It can therefore be excluded, that there is a $\mathcal{L}_\mu^{\mathrm{lin}}$-formula which defines the same property over $w_{2,1}$ but is finitely converging.

Chapter 6

Implications and Outlook

In Chap. 4 it was proven that all alternation-free $\varphi \in \mathcal{L}_\mu^{\mathrm{lin}}$ are finitely converging over all $w_{n,m}$. Additionally, a way how to use this result to show the finite convergence for arbitrary $\mathcal{L}_\mu^{\mathrm{lin}}$-formula over all $w_{n,m}$ was sketched. In contrast to this, in Chap. 5 it was shown that there are simple $\mathrm{HFL}^1_{\mathrm{lin}}$-formulas and $\mathrm{HFL}^2_{\mathrm{lin}}$-formulas which are not finitely converging over at least one instance of $w_{n,m}$. Furthermore, it was discussed that the $\mathrm{HFL}^1_{\mathrm{lin}}$-example defines a non-trivial property which is not definable by a $\mathcal{L}_\mu^{\mathrm{lin}}$-formula. With the understanding from Sec. 2.5 that the usage of $\mathcal{L}_\mu^{\mathrm{lin}}$ and $\mathrm{HFL}_{\mathrm{lin}}$ is no restriction to the validity of these findings for $\mathcal{L}_\mu$ and HFL, the results of this work can be summarized as follows: There are classes of structures that allow non-finite convergence of meaningful higher-order HFL-formulas in contrast to a guaranteed finite convergence of all $\mathcal{L}_\mu$-formulas. In addition to this, there are properties of these structures which are definable by a HFL-formula with non-finite convergence but not definable by any $\mathcal{L}_\mu$-formula. This implies that the higher expressive power of HFL compared to $\mathcal{L}_\mu$ also comes with higher or non-finite convergence.

In [BLL14] some techniques to eliminate fixpoint alternation in low-order fragments of HFL are presented. These techniques require that the HFL-formulas are considered in context of a special class of LTS. These special LTS are transition systems for which the considered (fixpoint) formula converges after a finite number of iterations. This classification leads to the following hierarchy of LTS: Let $\mathbb{T}^i_{\mathrm{fin}}$ be the class of LTS such that for all $\mathcal{T} \in \mathbb{T}^i_{\mathrm{fin}}$ and all $\varphi \in \mathrm{HFL}^i$ it holds that $\varphi \in \mathrm{FC}(\mathcal{T})$. Then, the hierarchy is given as follows:

$$\mathbb{T}^0_{\mathrm{fin}} \supseteq \mathbb{T}^1_{\mathrm{fin}} \supseteq \cdots \supseteq \bigcap_{i \in \mathbb{N}} \mathbb{T}^i_{\mathrm{fin}}.$$

The relation $\mathbb{T}^i_{\mathrm{fin}} \supseteq \mathbb{T}^{i+1}_{\mathrm{fin}}$ is directly implied by $\mathrm{HFL}^i \subseteq \mathrm{HFL}^{i+1}$. The strictness of this hierarchy is an open question. The findings of this work lay a foundation for a possible answer: It is strongly conjectured that $\mathbb{T}^0_{\mathrm{fin}} \neq \mathbb{T}^1_{\mathrm{fin}}$, which implies that $\mathbb{T}^0_{\mathrm{fin}} \supsetneq \mathbb{T}^1_{\mathrm{fin}}$. The case $\mathbb{T}^0_{\mathrm{fin}} \supsetneq \mathbb{T}^2_{\mathrm{fin}}$ is also considered by this work, but is already implied by $\mathbb{T}^0_{\mathrm{fin}} \neq \mathbb{T}^1_{\mathrm{fin}}$. For the alternation-free fragments of $\mathcal{L}_\mu$ respectively HFL^0 and HFL^1 the separation of the classes $\mathbb{T}^0_{\mathrm{fin}}$ and $\mathbb{T}^1_{\mathrm{fin}}$ is proven and for the general case the result does seem to be an adjacent step.

Further Research

There are several next steps, which could be taken in continuation of this work:

- The finite convergence of arbitrary $\mathcal{L}_\mu^{\mathrm{lin}}$-formulas over $w_{n,m}$ needs a formal proof.

The alternation-free case is proven, but the general case was only sketched in Sec. 4.3. An idea how to approach this was given, but there are crucial parts which still need to be worked out in detail.

- The family $w_{n,m}$ should be investigated thoroughly. Its properties were used in this work, but it is left open to formally capture a property or a combination of properties which cause the findings of this work.

- The methods used in the proof of Chap. 4 are tailored for the specific case of $\mathcal{L}_\mu$-calculus or HFL^0 formulas. In particular, this means the utilized equivalence of $\mathcal{L}_\mu$ and APA, which allows the use of automata-theoretic arguments. In hindsight of differentiating higher-order fragments of HFL, a redesigned version of this proof, avoiding such specific steps, could carve out a more general understanding.

- A more distant step is further investigating the strictness of the hierarchy presented in the previous section. It is suspected that clearing the situation for $\mathbb{T}^1_{\text{fin}}$ and $\mathbb{T}^2_{\text{fin}}$ leads to a sophisticated conjecture about the strictness in general. The introduced family $w_{n,m}$ holds a good starting point for further insights into this problem.

- This work is located in the area of convergence studies of fixpoints, especially fixpoints definable in the modal μ-calculus. This research direction is at best in the early stages, which causes a lack of tools and methods. For example, taking existing results about semantic equivalence of automata, games and logics and extending them to also include arguments about approximative semantics offers an interesting set of tools for further research.

Bibliography

[AL13] Bahareh Afshari and Graham E. Leigh. “On closure ordinals for the modal mu-calculus”. In: *Computer Science Logic 2013 (CSL 2013), CSL 2013, September 2-5, 2013, Torino, Italy*. Ed. by Simona Ronchi Della Rocca. Vol. 23. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2013, pp. 30–44. DOI: 10.4230/LIPIcs.CSL.2013.30. URL: https://doi.org/10.4230/LIPIcs.CSL.2013.30.

[ALS07] Roland Axelsson, Martin Lange, and Rafal Somla. “The Complexity of Model Checking Higher-Order Fixpoint Logic”. In: *Logical Methods in Computer Science* 3.2 (2007). DOI: 10.2168/LMCS-3(2:7)2007. URL: https://doi.org/10.2168/LMCS-3(2:7)2007.

[BBW07] Patrick Blackburn, J. F. A. K. van Benthem, and Frank Wolter, eds. *Handbook of Modal Logic*. Vol. 3. Studies in logic and practical reasoning. North-Holland, 2007. ISBN: 978-0-444-51690-9. URL: https://www.sciencedirect.com/bookseries/studies-in-logic-and-practical-reasoning/vol/3/suppl/C.

[BFL15] Florian Bruse, Oliver Friedmann, and Martin Lange. “On guarded transformation in the modal μ-calculus”. In: *Log. J. IGPL* 23.2 (2015), pp. 194–216. DOI: 10.1093/jigpal/jzu030. URL: https://doi.org/10.1093/jigpal/jzu030.

[BKP86] Howard Barringer, Ruurd Kuiper, and Amir Pnueli. “A Really Abstract Concurrent Model and its Temporal Logic”. In: *Conference Record of the Thirteenth Annual ACM Symposium on Principles of Programming Languages, St. Petersburg Beach, Florida, USA, January 1986*. ACM Press, 1986, pp. 173–183. DOI: 10.1145/512644.512660. URL: https://doi.org/10.1145/512644.512660.

[BLL14] Florian Bruse, Martin Lange, and Etienne Lozes. “Collapses of fixpoint alternation hierarchies in low type-levels of higher-order fixpoint logic”. In: *Workshop on Programming and Reasoning on Infinite Structures*. 2014.

[BOW14] Achim Blumensath, Martin Otto, and Mark Weyer. “Decidability Results for the Boundedness Problem”. In: *Logical Methods in Computer Science* 10.3 (2014). DOI: 10.2168/LMCS-10(3:2)2014. URL: https://doi.org/10.2168/LMCS-10(3:2)2014.

[CL08] Thomas Colcombet and Christof Löding. “The Non-deterministic Mostowski Hierarchy and Distance-Parity Automata”. In: *Automata, Languages and Programming, 35th International Colloquium, ICALP 2008, Reykjavik, Iceland, July 7-11, 2008, Proceedings, Part II - Track B: Logic, Semantics, and Theory of Programming & Track C: Security and Cryptography Foundations*. Ed. by Luca Aceto et al. Vol. 5126. Lecture Notes in Computer Science. Springer, 2008, pp. 398–409. DOI: 10.1007/978-3-540-70583-3_33. URL: https://doi.org/10.1007/978-3-540-70583-3_33.

[Cza10] Marek Czarnecki. "How Fast Can the Fixpoints in Modal mu-Calculus Be Reached?" In: *7th Workshop on Fixed Points in Computer Science, FICS 2010, Brno, Czech Republic, August 21-22, 2010*. Ed. by Luigi Santocanale. Laboratoire d'Informatique Fondamentale de Marseille, 2010, pp. 35–39. URL: `https://hal.archives-ouvertes.fr/hal-00512377/document\#page=36`.

[Dam94] Mads Dam. "CTL* and ECTL* as Fragments of the Modal mu-Calculus". In: *Theor. Comput. Sci.* 126.1 (1994), pp. 77–96. DOI: `10.1016/0304-3975(94)90269-0`. URL: `https://doi.org/10.1016/0304-3975(94)90269-0`.

[EC80] E. Allen Emerson and Edmund M. Clarke. "Characterizing Correctness Properties of Parallel Programs Using Fixpoints". In: *Automata, Languages and Programming, 7th Colloquium, Noordweijkerhout, The Netherlands, July 14-18, 1980, Proceedings*. Ed. by J. W. de Bakker and Jan van Leeuwen. Vol. 85. Lecture Notes in Computer Science. Springer, 1980, pp. 169–181. DOI: `10.1007/3-540-10003-2\_69`. URL: `https://doi.org/10.1007/3-540-10003-2\_69`.

[EL86] E. Allen Emerson and Chin-Laung Lei. "Efficient Model Checking in Fragments of the Propositional Mu-Calculus (Extended Abstract)". In: *Proceedings of the Symposium on Logic in Computer Science (LICS '86), Cambridge, Massachusetts, USA, June 16-18, 1986*. IEEE Computer Society, 1986, pp. 267–278.

[GKL14] Julian Gutierrez, Felix Klaedtke, and Martin Lange. "The μ-calculus alternation hierarchy collapses over structures with restricted connectivity". In: *Theor. Comput. Sci.* 560 (2014), pp. 292–306. DOI: `10.1016/j.tcs.2014.03.027`. URL: `https://doi.org/10.1016/j.tcs.2014.03.027`.

[GS19] Maria João Gouveia and Luigi Santocanale. "$\aleph_1$ and the modal μ-calculus". In: vol. 15. 4. 2019. URL: `https://lmcs.episciences.org/5808`.

[HL11] Martin Hofmann and Martin Lange. *Automatentheorie und Logik*. eXamen.press. Springer, 2011. ISBN: 978-3-642-18089-7. DOI: `10.1007/978-3-642-18090-3`. URL: `https://doi.org/10.1007/978-3-642-18090-3`.

[JW96] David Janin and Igor Walukiewicz. "On the Expressive Completeness of the Propositional mu-Calculus with Respect to Monadic Second Order Logic". In: *CONCUR '96, Concurrency Theory, 7th International Conference, Pisa, Italy, August 26-29, 1996, Proceedings*. Ed. by Ugo Montanari and Vladimiro Sassone. Vol. 1119. Lecture Notes in Computer Science. Springer, 1996, pp. 263–277. DOI: `10.1007/3-540-61604-7\_60`. URL: `https://doi.org/10.1007/3-540-61604-7\_60`.

[Kai95] Roope Kaivola. "Axiomatising Linear Time Mu-calculus". In: *CONCUR '95: Concurrency Theory, 6th International Conference, Philadelphia, PA, USA, August 21-24, 1995, Proceedings*. Ed. by Insup Lee and Scott A. Smolka. Vol. 962. Lecture Notes in Computer Science. Springer, 1995, pp. 423–437. DOI: `10.1007/3-540-60218-6\_32`. URL: `https://doi.org/10.1007/3-540-60218-6\_32`.

[Koz83] Dexter Kozen. "Results on the Propositional mu-Calculus". In: *Theor. Comput. Sci.* 27 (1983), pp. 333–354. DOI: `10.1016/0304-3975(82)90125-6`. URL: `https://doi.org/10.1016/0304-3975(82)90125-6`.

[KVW00] Orna Kupferman, Moshe Y. Vardi, and Pierre Wolper. “An automata-theoretic approach to branching-time model checking”. In: *J. ACM* 47.2 (2000), pp. 312–360. DOI: 10.1145/333979.333987. URL: https://doi.org/10.1145/333979.333987.

[Mat02] Radu Mateescu. “Local Model-Checking of Modal Mu-Calculus on Acyclic Labeled Transition Systems”. In: *Tools and Algorithms for the Construction and Analysis of Systems, 8th International Conference, TACAS 2002, Held as Part of the Joint European Conference on Theory and Practice of Software, ETAPS 2002, Grenoble, France, April 8-12, 2002, Proceedings*. Ed. by Joost-Pieter Katoen and Perdita Stevens. Vol. 2280. Lecture Notes in Computer Science. Springer, 2002, pp. 281–295. DOI: 10.1007/3-540-46002-0_20. URL: https://doi.org/10.1007/3-540-46002-0_20.

[MV19] Gian Carlo Milanese and Yde Venema. “Closure Ordinals of the Two-Way Modal μ-Calculus”. In: *Logic, Language, Information, and Computation - 26th International Workshop, WoLLIC 2019, Utrecht, The Netherlands, July 2-5, 2019, Proceedings*. Ed. by Rosalie Iemhoff, Michael Moortgat, and Ruy J. G. B. de Queiroz. Vol. 11541. Lecture Notes in Computer Science. Springer, 2019, pp. 498–515. DOI: 10.1007/978-3-662-59533-6_30. URL: https://doi.org/10.1007/978-3-662-59533-6_30.

[Niw86] Damian Niwinski. “On Fixed-Point Clones (Extended Abstract)”. In: *Automata, Languages and Programming, 13th International Colloquium, ICALP86, Rennes, France, July 15-19, 1986, Proceedings*. Ed. by Laurent Kott. Vol. 226. Lecture Notes in Computer Science. Springer, 1986, pp. 464–473. DOI: 10.1007/3-540-16761-7_96. URL: https://doi.org/10.1007/3-540-16761-7_96.

[Ott99] Martin Otto. “Eliminating Recursion in the μ-Calculus”. In: *STACS 99, 16th Annual Symposium on Theoretical Aspects of Computer Science, Trier, Germany, March 4-6, 1999, Proceedings*. Ed. by Christoph Meinel and Sophie Tison. Vol. 1563. Lecture Notes in Computer Science. Springer, 1999, pp. 531–540. DOI: 10.1007/3-540-49116-3_50. URL: https://doi.org/10.1007/3-540-49116-3_50.

[Tar+55] Alfred Tarski et al. “A lattice-theoretical fixpoint theorem and its applications.” In: *Pacific journal of Mathematics* 5.2 (1955), pp. 285–309.

[VV04] Mahesh Viswanathan and Ramesh Viswanathan. “A Higher Order Modal Fixed Point Logic”. In: *CONCUR 2004 - Concurrency Theory, 15th International Conference, London, UK, August 31 - September 3, 2004, Proceedings*. Ed. by Philippa Gardner and Nobuko Yoshida. Vol. 3170. Lecture Notes in Computer Science. Springer, 2004, pp. 512–528. DOI: 10.1007/978-3-540-28644-8_33. URL: https://doi.org/10.1007/978-3-540-28644-8_33.

[Wal00] Igor Walukiewicz. “Completeness of Kozen’s Axiomatisation of the Propositional μ-Calculus”. In: *Inf. Comput.* 157.1-2 (2000), pp. 142–182. DOI: 10.1006/inco.1999.2836. URL: https://doi.org/10.1006/inco.1999.2836.

[Wil01] Thomas Wilke. “Alternating Tree Automata, Parity Games, and Modal m-Calculus.” In: *Bulletin of the Belgian Mathematical Society Simon Stevin* 8.2 (2001), p. 359.