Finite Convergence of μ-Calculus Fixpoints on Genuinely Infinite Structures

Florian Bruse
School of Electrical Engineering and Computer Science, University of Kassel, Germany

Marco Sälzer
School of Electrical Engineering and Computer Science, University of Kassel, Germany

Martin Lange
School of Electrical Engineering and Computer Science, University of Kassel, Germany

Abstract

The modal μ-calculus can only express bisimulation-invariant properties. It is a simple consequence of Kleene's Fixpoint Theorem that on structures with finite bisimulation quotients, the fixpoint iteration of any formula converges after finitely many steps. We show that the converse does not hold: we construct a word with an infinite bisimulation quotient that is locally regular so that the iteration for any fixpoint formula of the modal μ-calculus on it converges after finitely many steps. This entails decidability of μ-calculus model-checking over this word. We also show that the reason for the discrepancy between infinite bisimulation quotients and trans-finite fixpoint convergence lies in the fact that the μ-calculus can only express regular properties.

2012 ACM Subject Classification Theory of computation → Modal and temporal logics; Theory of computation → Higher order logic; Theory of computation → Logic and verification; Theory of computation → Automata over infinite objects

Keywords and phrases temporal logic, fixpoint iteration, bisimulation

Digital Object Identifier 10.4230/LIPIcs.MFCS.2021.24

1 Introduction

The modal μ-calculus $\mathcal{L}_\mu$, as it was introduced by Kozen [17], has become a de-facto standard yardstick amongst formal specification languages for programs. It is obtained in a principally simple way, namely by extending standard modal logic with extremal fixpoint quantifiers. Since most operators used in temporal logics can be characterised as least or greatest fixpoints, $\mathcal{L}_\mu$ can embed standard temporal logics like CTL, LTL and CTL* [10]. $\mathcal{L}_\mu$ is also, in a sense, the largest regular program specification logic as it is equi-expressive to the bisimulation-invariant fragment of Monadic Second-Order Logic [13]. Hence, studying its model-theoretic properties helps to answer questions after what can and cannot be formally expressed about programs in regular specification languages.

Fixpoint formulas in $\mathcal{L}_\mu$ denote sets of states in a labelled transition system (LTS), and Kleene's Fixpoint Theorem [16] can be used to approximate such fixpoints in a chain of sets: for instance, the semantics of some fixpoint definition $\mu X.\varphi(X)$ can be approximated from below by the sequence of sets X^i, where $X^0 = \emptyset$ and X^{i+1} is obtained as the semantics of $\varphi(X^i)$. For infinite structures, it is generally necessary to extend this sequence to trans-finite ordinal numbers. Over any LTS whose state space forms a set, this sequence must stabilise eventually at precisely the least fixpoint of φ.

Recent times have seen increased interest in the details of this process, regarding questions of when exactly it stabilises, and how this depends on the underlying structure and formula in question. Such questions are not simplified by fixpoint alternation – the ability to nest mutually dependent fixpoint definitions of different kinds. It is known that, over the class

46th International Symposium on Mathematical Foundations of Computer Science (MFCS 2021).
Editors: Filippo Bonchi and Simon J. Puglisi; Article No. 24; pp. 24:1–24:19
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

of all LTS, this is unavoidable in that alternation-free formulas do not capture all of $\mathcal{L}_\mu$'s expressiveness [5]. However, some classes of structures are known, for instance words, over which any $\mathcal{L}_\mu$ formula is equivalent to an alternation-free one [15].

The first notable result regarding the question of fixpoint stabilisation in $\mathcal{L}_\mu$ is that it is decidable whether an $\mathcal{L}_\mu$ formula is equivalent to a modal one [20]. Other and more recent research concerns the ordinals at which fixpoint iteration stabilises, in particular which ordinals are candidates for such a bound. This concerns the question whether, for a given formula $\mu X.\varphi$, there is some ordinal α such that fixpoint iteration $X^0, X^1, \dots$ stabilises after at most α steps over *any* LTS. Such a (minimal) α is called the *closure ordinal* of $\mu X.\varphi$. Czarnecki [9] shows that, for each ordinal $\alpha < \omega^2$, there is some $\mathcal{L}_\mu$ formula with closure ordinal α. Afshari and Leigh [1] show that, for alternation-free formulas, ω^2 is a tight upper bound for closure ordinal candidates. [12] shows that ω_1 is the closure ordinal of some $\mathcal{L}_\mu$ formula, and [19] shows that all ordinals below ω^ω are closure ordinals for some formula in the two-way μ-calculus, i.e. the extension of $\mathcal{L}_\mu$ by backwards modalities. The situation in the intuitionistic setting is studied in e.g. [11].

The present paper is concerned with a slightly different but related question: we ask for closure ordinals on particular structures, i.e. at which ordinal do the iteration processes of *all* $\mathcal{L}_\mu$ formulas stabilise? Similar problems have been investigated by Barwise and Moschovakis in the context of first-order logic, see e.g. [4]. The problem of finding closure ordinals of classes of structures relates to the previous problem, since the former closure ordinals obviously bound the latter over the given class. Hence, studying closure ordinals of classes of structures contributes to the understanding of closure ordinals of formulas.

For $\mathcal{L}_\mu$, the cardinality of the structure in question is an obvious upper bound for its closure ordinal. Hence, on finite structures, fixpoint iteration must necessarily stabilise at some finite bound that is uniform for all formulas. This simple observation extends to structures with a *finite bisimulation quotient*, as an immediate consequence of $\mathcal{L}_\mu$'s inability to distinguish bisimilar states. An interesting question arises as the converse of this: does an inherently infinite structure, i.e. one whose bisimulation quotient is infinite, allow formulas to have an infinite fixpoint iteration process? Put differently, are there structures with an infinite bisimulation quotient such that fixpoint iteration for *any* $\mathcal{L}_\mu$ formula converges after finitely many steps? It is tempting to equate having a finite bisimulation quotient with the finite convergence of all $\mathcal{L}_\mu$ fixpoints, yet the answer to the latter question is "yes".

We construct an LTS – in fact, an infinite word w_∞ – which has an infinite bisimulation quotient but all $\mathcal{L}_\mu$ formulas are equivalent to some finite approximation over it. Locally it seems to be regular, and no $\mathcal{L}_\mu$ formula can "see" the non-regular global pattern in it. Hence, $\mathcal{L}_\mu$ fixpoints cannot exploit this non-regularity in order to only stabilise after more than finitely many iteration steps. Local regularity means that w_∞ is self-similar: it is made of building blocks of increasing size, and some postfixes of it are w_∞ again if one maps suitable building blocks to certain symbols in the word's alphabet.

The result on finite convergence over w_∞ is obtained by first reducing the question for arbitrary $\mathcal{L}_\mu$ formulas to that of alternation-free ones. The aforementioned alternation-hierarchy collapse cannot simply be used off-the-shelf here as it makes no statement about the preservation of (in-)finiteness of closure ordinals. We then reduce the question to that for $\mathcal{L}_\mu$ formulas only containing fixpoints of one sort. We transform this into the analysis of runs of a very rudimentary fragment of alternating parity automata over w_∞, exploiting its self-similar structure. Decidability of $\mathcal{L}_\mu$ model checking over w_∞ follows as a corollary.

Finally, we show that the reason for the discrepancy between infinite bisimulation quotients and infinite $\mathcal{L}_\mu$ closure ordinals is to be found in the regularity of $\mathcal{L}_\mu$'s expressive power. We show that there are formulas of HFL – a natural higher-order extension of $\mathcal{L}_\mu$ [22] – that do not have finite convergence on w_∞. This is interesting because HFL can define the Kleene fixpoint iteration for $\mathcal{L}_\mu$ [8], and this can be used to disentangle fixpoint alternation, albeit at the cost of blow-up in formula size and type order, and the restriction to structures on which each $\mathcal{L}_\mu$ formula is equivalent to some finite approximation. Thus, the result of the paper at hand implies that this is not necessarily restricted a priori to structures with finite bisimulation quotients, as the aforementioned w_∞ is a counterexample.

The paper is organised as follows. In Sect. 2 we recall necessary preliminaries. In Sect. 3 we take a detailed look at fixpoint iteration for $\mathcal{L}_\mu$ formulas to make the notion of "finite convergence" formal. In Sect. 4 we develop an automata-theoretic criterion for a class of words to have finite convergence. We use this in Sect. 5 to prove finite convergence for w_∞. Sect. 6 contains the considerations on modal fixpoints of higher-order on w_∞. Sect. 7 concludes with remarks on further work.

2 Preliminaries

Words and languages. An *alphabet* is a finite, nonempty set Σ of *letters*, denoted by $a, b, \ldots$. A Σ-word is a finite or infinite sequence of letters. Infinite words are also called ω-words. The empty word is denoted by ε. We write $w = a_1 a_2 \cdots a_n$ for finite words, and $w = a_1 a_2 \cdots$ for infinite words. Given two finite words u and v with $v \neq \varepsilon$, then $u \cdot v^\omega$ is an infinite word. In both the finite and the infinite case, $w[i]$ denotes the i-th letter of w. Σ^* is the set of finite Σ-words, a subset of which is a Σ-*language* (of finite words). Languages of ω-words have no role in this paper. If the alphabet is clear from context, we simply speak of words and languages.

Words are a special case of labelled transition systems for which the notion of bisimilarity is well-known. We can simplify the definition for word structures and call two positions in a word *bisimilar* if they have the same postfix, i.e. if the rest of the word is the same from both positions. Obviously, an ω-word with two bisimilar positions is ultimately periodic. The bisimulation quotient of a word is the quotient w.r.t. bisimilarity. It is either the word itself (when all positions are mutually non-bisimilar), or a lasso-shaped representation of it (when at least two, and then necessarily all following pairs of positions are bisimilar).

The modal μ-calculus. We introduce the modal μ-calculus in its linear-time version only. Let Σ be an alphabet, let $\mathcal{X}$ be a set of fixpoint variables. The syntax of the linear-time μ-calculus in negation normal form, just $\mathcal{L}_\mu$ from now on, is given by the grammar

$$\varphi ::= a \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \bigcirc \varphi \mid X \mid \mu X.\, \varphi \mid \nu X.\, \varphi$$

where $X \in \mathcal{X}$ and $a \in \Sigma$. Other connectives such as $\mathtt{tt}, \mathtt{ff}, \rightarrow$ etc. are defined as usual. Note that negation is definable using De Morgan, $\neg a \equiv \bigvee_{b \neq a} b$ and duality between μ and ν.

The notion of a subformula is standard. The *size* $|\varphi|$ of a formula φ is the number of its distinct subformulas. Fixpoint quantifiers $\sigma \in \{\mu, \nu\}$ act as variable binders. The notion of free and bound occurrence, as well as that of a closed formula are as usual. In a formula $\sigma X.\varphi$, the subformula φ is the *defining* formula of X. A variable bound by μ is a *least-fixpoint variable*. It is a *greatest-fixpoint variable* if it is bound by ν.

We assume formulas to be *well-named* in the sense that each fixpoint variable is bound at most once. Clearly, any formula is equivalent to a well-named one via renaming of variables. In a well-named formula φ, there is a function fp_φ that maps each fixpoint variable to the defining formula of this variable. We drop the index if φ is clear from context. Well-namedness induces a partial order $<_{\mathrm{fp}}$ defined via $X <_{\mathrm{fp}} Y$ iff $\mathrm{fp}_\varphi(X)$ is a proper subformula of $\mathrm{fp}_\varphi(Y)$. Note that, if $X <_{\mathrm{fp}} Y$, then X has no free occurrences in $\mathrm{fp}_\varphi(Y)$.

We call a formula *unipolar* if it only contains one kind of fixpoint quantifiers. We call a formula *alternation-free* if no variable bound by a least fixpoint quantifier appears freely in the defining formula of a greatest fixpoint quantifier, and vice versa. In alternation-free formulas, the fixpoint variables can be partitioned into sets $\mathcal{X}_1, \dots, \mathcal{X}_k$ such that

- fixpoint variables do not appear freely in the defining formulas of variables from another set of the partition,
- all variables of one set have the same polarity, and
- for all $X \in \mathcal{X}_i$ and $Y \in \mathcal{X}_j$ with $i \neq j$: if $X <_{\text{fp}} Y$ then $j < i$.

A closed formula that contains no fixpoint quantifiers is in Basic Modal Logic (ML). The notion of *modal depth* $\text{md}(\varphi)$ of φ is defined as usual: $\text{md}(a) = 0$, $\text{md}(\psi_1 \vee \psi_2) = \text{md}(\psi_1 \wedge \psi_2) = \max\{\text{md}(\psi_1), \text{md}(\psi_2)\}$ and $\text{md}(\bigcirc \psi) = 1 + \text{md}(\psi)$.

Let $\eta\colon \mathcal{X} \to 2^{\mathbb{N}}$ be an *environment*. The semantics of an $\mathcal{L}_\mu$ formula on an ω-word w is a set of positions defined inductively as follows:

$$[\![a]\!]^w_\eta = \{i \in \mathbb{N} \mid w[i] = a\} \qquad [\![X]\!]^w_\eta = \eta(X)$$
$$[\![\varphi \wedge \psi]\!]^w_\eta = [\![\varphi]\!]^w_\eta \cap [\![\psi]\!]^w_\eta \qquad [\![\mu X.\, \varphi]\!]^w_\eta = \bigcap\{U \subseteq \mathbb{N} \mid [\![\varphi]\!]^w_{\eta[X \mapsto U]} \subseteq U\}$$
$$[\![\varphi \vee \psi]\!]^w_\eta = [\![\varphi]\!]^w_\eta \cup [\![\psi]\!]^w_\eta \qquad [\![\nu X.\, \varphi]\!]^w_\eta = \bigcup\{U \subseteq \mathbb{N} \mid U \subseteq [\![\varphi]\!]^w_{\eta[X \mapsto U]}\}$$
$$[\![\bigcirc \varphi]\!]^w_\eta = \{i \in \mathbb{N} \mid i + 1 \in [\![\varphi]\!]^w_\eta\}$$

We say that φ is *satisfiable* under η over w if $[\![\varphi]\!]^w_\eta \neq \emptyset$. It is valid, written $\models \varphi$, if $[\![\varphi]\!]^w_\eta = \mathbb{N}$ for all w and η.

Two closed $\mathcal{L}_\mu$ formulas φ and ψ are equivalent, written $\varphi \equiv \psi$, if they define the same set on all words. We write $\varphi \equiv_w \psi$ to denote that φ and ψ define the same set on w, and $\varphi \equiv_{\mathcal{C}} \psi$ for a class of words $\mathcal{C}$ if φ and ψ define the same set on all words in $\mathcal{C}$.

Trivial Automata. A *trivial* Σ-automaton (TrA) has the form (Q, δ, q_I, F, b) where

- Q is a finite nonempty set of states with initial state $q_I \in Q$ and final states $F \subseteq Q$,
- $\delta\colon (Q \setminus F) \times \Sigma \to Q$ is the transition function and
- $b \in \{0, 1\}$.

A *run* of $\mathcal{A}$ starting from some position i in some ω-word w is a finite or infinite sequence $q_0, q_1, \dots$ of states with $q_0 = q_I$ and $q_{j+1} = \delta(q_j, w_{j+i})$. Note that if such a sequence is finite, then the last state must necessarily be in F since δ is total on $Q \setminus F$. A run is *accepting* if it is finite and $b = 1$, or if it is infinite and $b = 0$. The set of positions in an ω-word defined by $\mathcal{A}$ is the set of positions from which it has an accepting run. Two TrA are equivalent if they define the same set on every infinite Σ-word.

We write $q \xrightarrow{v} q'$ for $q, q' \in Q, v \in \Sigma^*$ to denote that $\mathcal{A}$ will be in state q' after reading the finite word v starting from q. This includes the case in which $\mathcal{A}$ does not read the entirety of v since it stops beforehand. In this case, $q' \in F$.

Trivial automata are reminiscent of DFA but they operate on infinite words. They can be thought of as restricted parity automata with a single priority b. If $b = 1$ then no infinite run is accepting. Acceptance is typically still possible by hitting a state and alphabet symbol to which the transition function assigns $\mathtt{tt}$ (as a Boolean combination of states). By making acceptance explicit through final states instead, we can define these trivial automata to be deterministic which is useful in proofs later on.

▶ **Lemma 1.** *Any unipolar $\mathcal{L}_\mu$ formula is equivalent to a TrA.*

Proof. (Sketch) It is well-known that $\mathcal{L}_\mu$ formulas on words can be translated into alternating parity automata (APA) [23]. Unipolar formulas result in APA of a single priority: 1 for least, 0 for greatest fixpoints. Since the APA has only one priority, determinisation is possible through a double powerset construction, paying attention to states with no successors by either introducing a sink state, or by making them final. ◀

3 Unfolding of Fixpoint Formulas

Single fixpoints. We formalise the notion of finite fixpoint convergence. First consider the case of a formula with a single fixpoint. Let $\mu X.\varphi$ be a formula, η be an environment, and w be a word. Then φ defines a monotonic function $f\colon T \mapsto [\![\varphi]\!]^w_{\eta[X\mapsto T]}$. Approximations to the least fixpoint are defined via

$$T^0_X = \emptyset \;,\quad T^{i+1}_X = f(T^i_X) = [\![\varphi]\!]^w_{\eta[X\mapsto T^i_X]} \;,\quad T^\omega_X = \bigcup_{i\in\mathbb{N}} T^i_X.$$

Since we restrict ourselves to word structures with no branching, we do not need to consider approximations beyond ω, and by Kleene's Fixpoint Theorem we have $T^\omega_X = [\![\mu X.\,\varphi]\!]^w_\eta$.

Note that the T^i_X are definable in $\mathcal{L}_\mu$ by (renaming instances of) formulas φ^i via $\varphi^0 = \mathtt{ff}, \varphi^{i+1} = \varphi[\varphi^i/X]$ independently of w and η. We call φ^i the *ith unfolding* of $\mu X.\varphi$. If over w we have $T^i_X = T^\omega_X$ for some i then the fixpoint is equivalent to its ith unfolding, and the set defined by it can also be defined by a formula without the fixpoint. The definitions for greatest fixpoints are analogue, except that one starts with $\mathbb{N}$ instead of $\emptyset$, and with $\mathtt{tt}$ instead of $\mathtt{ff}$.

Multiple and nested fixpoints. For formulas containing more than one fixpoint subformula, possibly in a nested way, it is less clear what it means for these subformulas to be "unfolded n times", as there is some ambiguity w.r.t. the order in which the participating formulas are to be unfolded. The literature also contains no standard agreed-upon definition. Instead, one of the following three constructions is employed as the respective authors see fit: Unfolding bottom-up, unfolding top-down, or unfolding on demand, following a construction seen in e.g. [21]. We review all three of them here and show that they produce the same formula (cf. Lem. 7). Hence, using either of them as suitable is permissible, and the results obtained in Sect. 5 later on hold for any of these reasonable interpretations of finite fixpoint convergence.

This thorough discussion is necessitated by to two reasons: Note that we have very strict requirements with respect to unfolding procedures in the sense that, for an "nth unfolding", whenever a fixpoint is to be unfolded during the procedure, it is unfolded exactly n times. This differs from other notions where the amount of unfoldings can vary from fixpoint to fixpoint as long as some fixpoint-free formula is produced. Moreover, contrary to the case of only one fixpoint, monotonicity of the unfolding can be lost for formulas that are not unipolar, contradicting what one might intuitively assume (cf. Ex. 8), in particular w.r.t. stabilisation of the process.

The first unfolding procedure is quite straightforward: Pick a formula that is minimal w.r.t. $<_{\text{fp}}$ and unfold it n times. Clearly, this procedure terminates after k steps, if the formula in question contains k fixpoint definitions. However, it is not immediately clear whether a common n exists if one is interested in producing a formula equivalent to the original one over some structure.

▶ **Definition 2** (Bottom-up unfolding). *Define* $\hat{\mu} := \mathtt{ff}$ *and* $\hat{\nu} := \mathtt{tt}$*. Let* $n \geq 0$ *and let* φ *be an* $\mathcal{L}_\mu$ *formula. Let* $X_1, \ldots, X_k$ *be an enumeration of its fixpoint variables such that* X_i *is bound by some* $\sigma X_i.\psi_i$*, and* $X_i \not<_{\mathrm{fp}} X_j$ *for* $j > i$*. Let* $\varphi_k^n, \ldots, \varphi_0^n$ *with* $\varphi_k^n = \varphi$ *be a sequence of formulas defined via*

$$\psi_i^0 = \hat{\sigma_i} \ , \qquad \psi_i^{j+1} = \mathrm{fp}_{\varphi_i^n}(X_i)[\psi_i^j/X_i] \ , \qquad \varphi_{i-1}^n = \varphi_i^n[\psi_i^n/\sigma_i X_i.\, \mathrm{fp}_{\varphi_i^n}(X_i)]$$

Then φ_0^n *is the* nth bottom-up unfolding *of* φ*.*

▶ **Example 3.** Let $\varphi = \nu X. \bigcirc(\mu Y.X \wedge Y)$ and let $n = 2$. Clearly $Y <_{\mathrm{fp}} X$. Hence, $\psi_2^2 = X \wedge (X \wedge \mathtt{ff})$ whence $\varphi_1^2 = \nu X. \bigcirc(X \wedge (X \wedge \mathtt{ff}))$. Moreover $\psi_1^1 = \bigcirc(\mathtt{tt} \wedge (\mathtt{tt} \wedge \mathtt{ff}))$ and, hence, $\varphi_0^2 = \bigcirc\big(\bigcirc(\mathtt{tt} \wedge (\mathtt{tt} \wedge \mathtt{ff})) \wedge (\bigcirc(\mathtt{tt} \wedge (\mathtt{tt} \wedge \mathtt{ff})) \wedge \mathtt{ff})\big)$.

The second procedure is perhaps the most straightforward one: Given some formula that contains fixpoints, and some n, pick some fixpoint definition that is maximal w.r.t $<_{\mathrm{fp}}$ and unfold it n times. Given that this may duplicate fixpoint definitions that are smaller w.r.t. $<_{\mathrm{fp}}$, this requires renaming and raises questions regarding termination of the procedure. Moreover, if one is interested into producing equivalent formulas over some structure, it is, again, not immediately clear whether a common n exists that can be used for all fixpoint definitions simultaneously. We start with an example.

▶ **Example 4.** Consider again the formula $\nu X. \bigcirc(\mu Y.X \wedge Y)$. Let $\psi = \bigcirc(\mu Y.X \wedge Y)$. By unfolding X twice as per above, we obtain the sequence of formulas $\psi^0 = \mathtt{tt}, \psi^1 = \bigcirc(\mu Y.\mathtt{tt} \wedge Y), \psi^2 = \bigcirc(\mu Y.(\bigcirc(\mu Y.\mathtt{tt} \wedge Y)) \wedge Y)$. Not only are there now two fixpoint definitions involving Y, their variables are also comparable via $<_{\mathrm{fp}}$. However, by renaming one of them, we obtain the formula $\bigcirc(\mu Y.(\bigcirc(\mu Y'.\mathtt{tt} \wedge Y')) \wedge Y)$. Note that the two variables are not mutually recursive since Y does not appear in the defining formula of Y'.

Since the two variables are comparable, and $Y >_{\mathrm{fp}}> Y'$, we proceed by unfolding it which, after renaming, results in $\bigcirc((\bigcirc(\mu Y'.\mathtt{tt} \wedge Y')) \wedge ((\bigcirc(\mu Y''.\mathtt{tt} \wedge Y'')) \wedge \mathtt{ff}))$.

Again, we obtain two fixpoint definitions. However, this time, the two variables in question are not comparable via $<_{\mathrm{fp}}$, whence the order of unfolding is obviously not important. We unfold Y' first and obtain $\bigcirc((\bigcirc(\mathtt{tt} \wedge (\mathtt{tt} \wedge \mathtt{ff}))) \wedge ((\bigcirc(\mu Y''.\,\mathtt{tt} \wedge Y'')) \wedge \mathtt{ff}))$ and then, after unfolding Y'', the formula $\bigcirc((\bigcirc(\mathtt{tt} \wedge (\mathtt{tt} \wedge \mathtt{ff}))) \wedge ((\bigcirc(\mathtt{tt} \wedge (\mathtt{tt} \wedge \mathtt{ff}))) \wedge \mathtt{ff}))$.

This is the same formula as the one obtained by bottom-up unfolding in Ex. 3. This is in fact no coincidence, cf. Lemma 7 below.

▶ **Definition 5** (Top-down unfolding). *Let* $n \geq 0$ *and let* φ *be an* $\mathcal{L}_\mu$ *formula. Define a sequence* $\varphi_0^n, \varphi_1^n, \ldots$ *where* $\varphi_0^n = \varphi$*, and* φ_{i+1}^n *is obtained from* φ_i^n *via the following process: if* φ_i^n *contains no fixpoint definitions,* $\varphi_{i+1}^n = \varphi_i^n$*. Otherwise, let* X *be a variable that is maximal w.r.t.* $<_{\mathrm{fp}}$ *in* φ_i^n*. Let* $\sigma X.\,\psi$ *be the subformula that defines* X*. Define* $\psi^0 = \hat{\sigma}$ *and* $\psi^{j+1} = \psi[\psi^j/X]$*. Then* $\varphi_{i+1}^n = \varphi_i^n[\psi'^n/\sigma X.\,\psi]$ *where* ψ'^n *is a copy of* ψ^n *made well-named via renaming of variables. If* $\varphi_i^n = \varphi_{i+1}^n$*, then the* nth top-down unfolding *of* φ *is* φ_i^n*.*

As already said, is not immediately obvious that the above process terminates, but Ex. 4 already gives a hint. Unfolding a fixpoint formula may duplicate other, inner fixpoint formulas but the duplicates are independent of each other. Unfolding the outer may create further duplicates of inner duplicates, but these are not mutually recursive, which gives a termination argument.

In order to not deal with ambiguities around the termination of the process, and to avoid issues around unfolding formulas containing free fixpoint variables (cf. the bottom-up approach), we review a third definition of the nth unfolding of a fixpoint formula centered around tracking for each fixpoint how often it has been unfolded already. This procedure is also folklore and based on the well-known notion of μ-signatures [21] or techniques used to unfold parity automata into $\mathcal{L}_\mu$-formulas [7].

▶ **Definition 6.** *Let $n \geq 0$ and let φ be an $\mathcal{L}_\mu$ formula. Let $X_1, \dots, X_k$ be an enumeration of its fixpoint variables such that $X_i \not<_{\mathrm{fp}} X_j$ for $j > i$ and such that σ_i denotes the (polarity of the) fixpoint quantifier for X_i. For a tuple $s = (c_1, \dots, c_k)$ let $s(i) = c_i$ if $1 \leq i \leq k$ and, if $s(i) > 0$, define $s[i{-}{-}]$ as the k-tuple $(c_1, \dots, c_i - 1, n, \dots, n)$.*

Define φ^n as φ^{s_I}, where $s_I = (n, \dots, n)$ and ψ^s is given inductively as

$$a^s = a \qquad (\psi_1 \vee \psi_2)^s = \psi_1^s \vee \psi_2^s$$
$$(\neg\psi)^s = \neg\psi^s \qquad (\psi_1 \wedge \psi_2)^s = \psi_1^s \wedge \psi_2^s$$
$$(\bigcirc \psi)^s = \bigcirc \psi^s \qquad X_i^s = (\sigma X_i.\, fp_\varphi(X_i))^s = \begin{cases} \hat{\sigma}_i & , \text{ if } s[i] = 0 \\ \mathrm{fp}_\varphi(X_i)^{s[i--]} & , \text{ otherwise.} \end{cases}$$

Clearly, φ^{s_I} is well-defined and fixpoint-free. Well-definedness follows from the fact that $s[i{-}{-}]$ is smaller than s in the lexicographical ordering. Note that we do not have to deal with well-namedness since no intermediate formulas containing fixpoint definitions occur due to the inductive definition centered around s.

We now establish that all three definitions given above actually produce the same formulas:

▶ **Lemma 7.** *Let φ be an $\mathcal{L}_\mu$ formula. Then the bottom-up unfolding of φ (cf. Def. 2) and the top-down unfolding of φ (cf. Def. 5) are equivalent to the unfolding defined in Def. 6. In particular, the top-down unfolding is well-defined.*

The proof has been moved to the appendix. It mostly consists of tracking the various substitutions.

We say that φ is equivalent to its nth unfolding over some word w if $\varphi \equiv_w \varphi^m$ for all $m \geq n$, i.e. if φ^n defines the same set on each of these words, and so do all further unfoldings. Note that, contrary to the case of a single fixpoint variable, it is not automatically the case that if $\varphi^n \equiv_w \varphi$, then $\varphi^{n+1} \equiv_w \varphi$. To illustrate this, consider the following example:

▶ **Example 8.** Let $\varphi = \nu X.\mu Y.(a \wedge \bigcirc X) \vee \bigcirc Y$. It defines the set of all positions after which a occurs infinitely often. Its first unfoldings are

$$\begin{aligned}
\varphi^0 &= \mathtt{tt} \\
\varphi^1 &= (a \wedge \bigcirc \mathtt{tt}) \vee \bigcirc \mathtt{ff} \equiv a \\
\varphi^2 &= (a \wedge \bigcirc((a \wedge \bigcirc \mathtt{tt}) \vee \bigcirc((a \wedge \bigcirc \mathtt{tt}) \vee \bigcirc \mathtt{ff}))) \\
&\quad \vee \bigcirc((a \wedge \bigcirc((a \wedge \bigcirc \mathtt{tt}) \vee \bigcirc((a \wedge \bigcirc \mathtt{tt}) \vee \bigcirc \mathtt{ff}))) \vee \bigcirc \mathtt{ff}) \\
&\equiv (a \wedge \bigcirc(a \vee \bigcirc a)) \vee \bigcirc((a \wedge \bigcirc(a \vee \bigcirc a)))
\end{aligned}$$

Take $w = (ba)^\omega$. Then φ^0 obviously defines $\mathbb{N}$, while φ^1 defines $\{2n+1 \mid n \in \mathbb{N}\}$ and then φ^2 again defines $\mathbb{N}$ and so do all further approximations. Similar examples can be constructed to separate any two approximations. In fact, over a word of the form $b^1ab^2ab^3a\cdots$, the formula φ is not equivalent to any of its unfoldings φ^i with $i \geq 1$, but still defines $\mathbb{N}$.

Note that φ from Ex. 8 is not unipolar. For unipolar formulas, monotonicity can be used to show that $\models \varphi^i \rightarrow \varphi^{i+1}$ for least fixpoint formulas, resp. $\models \varphi^{i+1} \rightarrow \varphi^i$ for greatest fixpoint formulas holds for all i.

4 Finite Fixpoint Convergence for $\mathcal{L}_\mu$

In this section we define the notion of a word having finite convergence, i.e. the property that all formulas, or all formulas of a certain kind, are equivalent to a finite unfolding over this word. We also develop a sufficient criterion in terms of runs of TrA, for this to hold.

▶ **Definition 9.** *Let w be an infinite word and let Φ be a set of $\mathcal{L}_\mu$ formulas. We say that w has finite convergence for Φ if, for every $\varphi \in \Phi$, there is n such that φ is equivalent to φ^n over w. We say that w has finite convergence for $\mathcal{L}_\mu$, if the above holds for the set of all $\mathcal{L}_\mu$ formulas.*

The rest of the section is devoted to reducing finite convergence over w for the set of $\mathcal{L}_\mu$ formulas to a rather simple criterion on the runs of TrA over w. Lemmas 10 and 11 establish that, if a word has finite convergence for the set of alternation-free formulas, it also has finite convergence for the set of all $\mathcal{L}_\mu$ formulas. The rest of the section establishes a criterion for a word to have finite convergence for the set of alternation-free formulas.

▶ **Lemma 10.** *Let w be an infinite word that has finite convergence for the set of all alternation-free $\mathcal{L}_\mu$ formulas. Then, every closed $\mathcal{L}_\mu$ formula φ of the form $\sigma X.\psi$ is equivalent over w to one in* ML*, and so are all its approximations X^i. Moreover, there is some i such that φ agrees with X^i over w.*

Proof. Let w and $\varphi = \mu X.\psi$. The case of $\sigma = \nu$ is analogous. By [15], φ is equivalent to an alternation-free formula over the class of all words. Then, by the assumption of the lemma, there is $\varphi' \in$ ML that is equivalent to this alternation-free formula, obtained via some finite unfolding of φ. Now let X^i be the ith approximation of φ and let ψ_i be the $\mathcal{L}_\mu$ formula that defines it. Note that it possibly contains fixpoint definitions, since the only fixpoint to be unfolded is X. However, with the same argument as before we obtain that ψ_i also must be equivalent to some alternation-free formula and, hence, to some formula ψ'_i in ML.

Regarding the claim that one of the approximations is already equivalent to φ, assume that this is not the case. Since the ψ'_i are obtained as formulas equivalent to finite approximations of φ, we must have that for each $i \in \mathbb{N}$, there must be some position $j_i \in [\![\psi'_{i+1}]\!]^w \setminus [\![\psi'_i]\!]^w$ and, hence, $j_i \in [\![\varphi']\!]^w \setminus [\![\psi'_i]\!]^w$. Consider the set $\Phi = \{\varphi'\} \cup \{\neg\psi'_i \mid i \in \mathbb{N}\}$. We show that it is satisfiable using the Compactness Theorem for ML. Consider any finite subset Ψ of Φ, w.l.o.g. it is of the form $\{\varphi'\} \cup \{\neg\psi'_i \mid i \leq k\}$ for some k. By the above, Ψ is satisfiable by a postfix of w, starting at j_k. Hence, Φ is also satisfiable, i.e. there is an ω-word w' that satisfies φ', but none of the ψ'_i. This is a contradiction, since $[\![\varphi']\!]^{w'} = \bigcup_{i\in\mathbb{N}}[\![\psi'_i]\!]^{w'}$ by definition. This contradiction stems from the assumption that there is not already some i such that $\psi'_i \equiv_w \varphi' \equiv_w \varphi$. This finishes the proof. ◀

▶ **Lemma 11.** *Let w be an infinite word. If w has finite convergence for the set of all alternation-free $\mathcal{L}_\mu$ formulas, it has finite convergence for the set of all $\mathcal{L}_\mu$ formulas.*

Proof. Let φ be an $\mathcal{L}_\mu$ formula. Using Lem. 10, we can obtain a non-uniform unfolding φ' of φ that is equivalent to φ over w, i.e. we show that there is m such that, following the pattern of the top-down unfolding procedure in Def. 5, for each fixpoint subformula there is some $n \leq m$ such that unfolding it n times yields an equivalent subformula. In a second step, we show that we also obtain a formula equivalent to φ if we unfold all fixpoint subformulas exactly m times. This is not immediately obvious due to the non-monotonicity seen in Ex. 8.

Towards the first goal, define a sequence $\varphi_0, \varphi_1, \ldots$ where $\varphi_0 = \varphi$, and φ_{i+1} is obtained from φ_i via the following process, similarly to the top-down unfolding: if φ_i contains no fixpoint definitions $\varphi_{i+1} = \varphi_i$. Otherwise, let X be a variable in φ_i that is maximal w.r.t

$<_{\text{fp}}$. Let $\sigma X.\,\psi$ be the subformula that defines X. By Lem. 10, there is some m_i such that the m_ith unfolding of ψ is equivalent to $\sigma X.\,\psi$, defined via $\psi^0 = \hat{\sigma}$ and $\psi^{i+1} = \psi[\psi^i/X]$. Then $\varphi'_{i+1} = \varphi_i[\psi^{m_i+1}/\sigma X.\,\psi]$ resulting from the (m_i+1)th unfolding of X. Let φ_{i+1} be a obtained from φ'_{i+1} via renaming such that φ_{i+1} is well-named. Note that the amount of times each fixpoint is unfolded varies. This is where the above differs from the top-down unfolding. The above process stabilises for the same reason the top-down unfolding is well-defined: unfolding an outermost fixpoint formula will create closed formulas, i.e. the ψ^i as described above are all closed. Hence, while unfolding an outermost fixpoint X can duplicate fixpoints smaller than X w.r.t $<_{\text{fp}}$, the defining formulas of the duplicates reside in different instances of the ψ^i and, hence, are not mutually recursive. In particular, unfolding a formula in ψ^{i+1} may create further duplicates by replicating ψ^i, but since ψ^i is closed, these further duplicates can then be unfolded independently of each other.

Hence, let φ' be the formula that results once this process stabilises. Note that, however, the unfolding is not necessarily uniform. Let $m = 1 + \max\{m_i \mid i \in \mathbb{N}\}$. Since the process stabilises, this maximum exists. We claim that φ is equivalent to φ^m over w. We show this by unfolding φ using the top-down procedure. Note that above, we have established that, for each fixpoint X in the process, there is some $m_i < m$ such that unfolding it $m_i + 1$ times, once it is X's turn, results in a formula equivalent to the one before. We now show that this property is kept if we instead unfold it m times. Let X be such a variable, and assume that the property holds so far. Note that X must be outermost by now. Let $\sigma X.\,\psi$ be the defining formula of X. We compare ψ^{m_i} and ψ^m, which are equivalent due to Lem. 10. Note that ψ^{m_i+1} and ψ^{m_i} are equivalent due to the definition of m_i. Since $m \geq m_i$, we have that $\psi^m = \psi^{m_i+k}$ for some k, and it contains ψ^{m_i} as a subformula. Since that subformula is closed, clearly the invariant holds for all fixpoint definitions in ψ^{m_i}, since the process inside this subformula will play out exactly like before. If $m_i + 1 = m$, we are done. Otherwise, consider a subformula in $\psi^{m_i+1+k'}$ for some $k' \leq m - m_i - 1$, but not in ψ^{m_i+1}, i.e., it is in the extra part of the formula due to the extra unfolding. Since $\psi^{m_i+1} \equiv_w \psi^{m_i}$ by definition, we also have that $\psi^{m_i+k'} \equiv_w \psi^{m_i}$. In other words, the part of the formula where X used to be, but some ψ^j has been substituted, is equivalent over w due to the definition of m_i. Moreover, both substituted formulas are closed and, hence, can be exchanged without interfering with unfolding of fixpoint formulas. Hence, for the purposes of fixpoint unfolding, all fixpoint formulas in ψ^m, but not in ψ^{m_i+1} behave like a fixpoint formula in ψ^{m_i+1}, but not in ψ^{m_i}. Since the invariant holds for the latter, it must also hold for the former.

It follows that we can make the unfolding uniform by just using the top-down unfolding process with m. Moreover, any $m' \geq m$ yields the same result by the same reasoning. This finishes the proof. ◀

▶ Remark 12. Note that Lem. 11 yields more than just the collapse of $\mathcal{L}_\mu$ to ML over w which can already be inferred from the collapse of $\mathcal{L}_\mu$ to alternation-free $\mathcal{L}_\mu$ over the class of all words (see [15]). Lem. 11 yields that every $\mathcal{L}_\mu$ formula φ is equivalent, over w, not only to some ML formula, but one obtained as an unfolding of φ (cf. also the remarks after Def. 9).

▶ **Lemma 13.** *Let w be an infinite word. If w has finite convergence for the set of all unipolar $\mathcal{L}_\mu$ formulas, it has finite convergence for the set of all alternation-free $\mathcal{L}_\mu$ formulas.*

The proof is a standard induction on the alternation classes using the fact that alternation-free $\mathcal{L}_\mu$ is obtained by capture-avoiding substitution of unipolar formulas. It is spelled out in the appendix.

▶ **Definition 14.** *Let w be an ω-word and $\mathcal{A}$ be a TrA. We say that $\mathcal{A}$ has k-*bounded runs *on w if all finite runs of $\mathcal{A}$ in w are of length k or less. We say that TrA have bounded runs on w if, for every TrA $\mathcal{A}$, there is k such that $\mathcal{A}$ has k-bounded runs on w.*

Clearly, if $\mathcal{A}$ has k-bounded runs on w, then acceptance of $\mathcal{A}$ can be expressed by a ML-formula of modal depth at most k, i.e., there is some ML-formula φ of modal depth at most k such that $[\![\varphi]\!]^w = \{i \mid \mathcal{A} \text{ accepts from } i\}$.

▶ **Lemma 15.** *Let w be an infinite word. If TrA have bounded runs over w then w has finite convergence for all unipolar formulas.*

Proof. We only show the result for least fixpoint formulas; for greatest it is analogous. Let w be given, φ be unipolar containing only least fixpoints, and φ^i be its ith unfolding. Since φ is unipolar, it is equivalent to a TrA $\mathcal{A}$ by Lem. 1. By the assumption, there is k such that if $\mathcal{A}$ halts on w from some position then it halts in k steps or less. Hence, acceptance of $\mathcal{A}$ on w can be expressed by some ML formula ψ, which means that ψ is equivalent to φ over w. Then, similar to the proof of Lem. 10, we can use the Compactness Theorem to obtain that φ is already equivalent to some φ^i. ◀

Lemmas 11, 13 and 15 yield the following.

▶ **Corollary 16.** *Let w be an infinite word. If TrA have bounded runs on w, then w has finite convergence for $\mathcal{L}_\mu$.*

5 A Word with Finite Convergence

We are now ready to give the construction of a word with finite fixpoint convergence. Let $\Sigma = \{a, b\}$. We define w_∞ using the following mutually recursive definitions of families of finite words α_i, β_i:

$$\alpha_0 = a \ , \quad \beta_0 = b \ , \quad \alpha_{i+1} = \alpha_i\alpha_i\beta_i\alpha_i\alpha_i \ , \quad \beta_{i+1} = \beta_i\beta_i\alpha_i\beta_i\beta_i$$

For example, $\alpha_2 = aabaa\ aabaa\ bbabb\ aabaa\ aabaa$. Then $w_\infty = \alpha_0\alpha_1\cdots$.

We obtain the following properties of w_∞:

▶ **Lemma 17.** *Let w_∞ be as above. Then it holds that*

1. *the length of α_i and β_i is 5^i,*
2. *α_i and β_i do not overlap, i.e. the minimal size for a word that contains both of them is $2 \cdot 5^i$,*
3. *the postfix of w_∞ starting after position $\sum_{j=0}^{i-1} 5^j$, i.e. at the first occurrence of α_i, can be considered a word in $\{\alpha_i, \beta_i\}^\omega$,*
4. *α_i, respectively β_i occurs at most 4 times in a row before the other one occurs.*
5. *the distance from any position to the next occurrence of one of α_i or β_i is at most $5^i - 1$,*
6. *The postfix of w_∞ starting after position $\sum_{j=0}^{i-1} 5^j$ is $h(w_\infty)$ for the homomorphism h with $h(a) = \alpha_i$ and $h(b) = \beta_i$,*
7. *w_∞ has an infinite bisimulation quotient.*

Proof. Item 1 is an immediate consequence of the definition of w_∞. Item 2 follows from a straightforward induction: The claim is obvious for $i = 0$, and for $i > 0$ we note that any potential overlap of α_{i+1} and β_{i+1} either induces overlap of α_i and β_i, or contradicts the fact that α_{i+1} contains only one occurrence of β_i. Item 3 follows from the construction of α_{i+1}, resp. β_{i+1}. Towards Item 4, we use Item 2 to note that consecutive occurrences of α_i, resp.

β_i must occur aligned to the building pattern of w_∞, i.e. following the pattern exhibited in Item 3 applied to α_{i+1} and β_{i+1}. The claim then follows directly from the construction of α_{i+1} and β_{i+1}.

Regarding Item 5, note that, by Item 3, eventually, w_∞ consists entirely of a sequence of α_i and β_i. Hence, one must occur after a distance of at most $5^i - 1$. Moreover, since the first occurrence of α_i is at position $\sum_{j=0}^{i-1} 5^j \leq 5^i - 1$, the claim also holds for the initial part of the word. Item 6 follows from Item 3 and the building pattern of w_∞.

It remains to prove Item 7, i.e. that w_∞ has an infinite bisimulation quotient. This holds since all positions of the form $\sum_{j=0}^{i-1} 5^j$, i.e. the first occurrences of α_i for $i \in \mathbb{N}$, are pairwise not bisimilar. Towards this, note that α_i^3 is the word following at position $\sum_{j=0}^{i-1} 5^j$, since the α_{i+1} following the first occurrence of α_i begins with α_i^2. Conversely, all positions of the form $\sum_{j=0}^{i'-1} 5^j$ with $i' > i$ mark the beginning of the first $\alpha_{i'}$, which begins with $\alpha_i^2\beta_i$ by construction. Hence, the positions $\sum_{j=0}^{i-1} 5^j$ and $\sum_{j=0}^{i'-1} 5^j$ for $i' > i$ are not bisimilar, which yields infinitely many pairwise not bisimilar positions. ◀

The aim now is to show that TrA have finite runs on w_∞. Let $\mathcal{A} = (Q, \delta, q_I, F, b)$ be a TrA, fixed for the remainder of the section. W.l.o.g. $b = 1$ for the remainder of the section, the proof for $b = 0$ is completely symmetric. Consider the subsets $A_0, A_1, \ldots \subseteq Q$ and $B_0, B_1, \ldots \subseteq Q$ defined via $q \in A_i$ iff $q \xrightarrow{\alpha_i} q'$ for some $q' \in F$ and $q \in B_i$ iff $q \xrightarrow{\beta_i} q'$ for some $q' \in F$.

Clearly, $A_i \subseteq A_{i+1}$ for all $i \geq 0$ since α_{i+1} starts with α_i, whence any word that begins with α_{i+1} also begins with α_i. Moreover, since $|Q| < \infty$, there must be $i, h \in \mathbb{N}$ such that $A_j = A_i$ for all $j \geq i$ and $B_j = B_h$ for all $j \geq h$. Let $k = 1 + \max\{i, h\}$, $A = A_k$ and $B = B_k$. Note that $A \cap B$ can be nonempty, and both A and B can be empty. Let $M = Q \setminus (A \cup B)$. Then M is the set of states such that $\mathcal{A}$ will not have accepted if it reads α_j or β_j for any j.

We now show that the self-similarity (cf. Lem. 17.6) and w_∞ eventually becoming almost featureless from the perspective of a bounded-memory automaton (cf. Lem. 17.3), together imply that a TrA can get trapped in M if it does not escape it fast enough.

▶ **Lemma 18.** *Let $q \in M$ and $j \geq k$. Then $q \xrightarrow{\alpha_j} q'$ for some $q' \in M$, and $q \xrightarrow{\beta_j} q'$ for some $q' \in M$.*

Proof. Let $q_0 \in M$, $j \geq k$. Remember that $\alpha_j = \alpha_{j-1}^2\beta_{j-1}\alpha_{j-1}^2$. Note that $q_0 \xrightarrow{\alpha_{j-1}} q_1$ for some $q_1 \notin A$, because otherwise $q_0 \xrightarrow{\alpha_j} q_2$ for some $q_2 \in F$ contradicting $q_0 \in M$. Moreover, $q_0 \xrightarrow{\alpha_{j-1}\alpha_{j-1}} q_2$ for some $q_2 \notin A$. If it were the case that $q_2 \in A$, then there would be q_1 with $q_0 \xrightarrow{\alpha_{j-1}} q_1$ and $q_1 \xrightarrow{\alpha_{j-1}} q_2$. Since $q_2 \in A$, there must be $q_3 \in F$ such that $q_2 \xrightarrow{\alpha_{j-1}} q_3$. Hence, $q_1 \xrightarrow{\alpha_j} q_3$, which implies $q_1 \in A$. This contradicts the previous result, whence $q_2 \notin A$.

In summary, $q \xrightarrow{\alpha_{j-1}} q'$ or $q \xrightarrow{\alpha_{j-1}\alpha_{j-1}} q'$ for a state $q \in M$ implies $q' \notin A$ and it can be easily inferred that $q \xrightarrow{\alpha_{j-1}\alpha_{j-1}} q'$ also implies $q' \notin B$. The same holds symmetrically for β_{j-1}. Now, these findings can be used to prove the lemma. Per assumption, the automaton does not halt reading α_j starting from q_0. Hence, there are q_1, q_2, q_3 such that $q_0 \xrightarrow{\alpha_{j-1}\alpha_{j-1}} q_1, q_1 \xrightarrow{\beta_{j-1}} q_2$ and $q_2 \xrightarrow{\alpha_{j-1}\alpha_{j-1}} q_3$. From the findings above it immediately follows that $q_1 \in M$. With the symmetric arguments for β_j it follows that $q_2 \notin B$ and from the fact that $q_0 \in M$ it also follows that $q_2 \notin A$, whence $q_2 \in M$. Then, with the same arguments as for q_1 and, we obtain that $q_3 \in M$, too. The case for β_j is analogous. ◀

We are now ready to prove that all TrA are bounded over the word w.

▶ **Theorem 19.** *TrA have bounded runs on w_∞.*

Proof. Let $\mathcal{A}$ be a TrA, M, k be as above and let l be some position in w_∞. We show that if $\mathcal{A}$ accepts from position l, then it does so within $6 \cdot 5^k - 1$ many steps. Let l' be the first position after l from which α_k or β_k starts. Let u be the word from position l to position l'. If $q_I \xrightarrow{u} q$ for $q \in F$, we are done since by Lem. 17.5 we have $|u| \leq 5^k - 1$.

Otherwise, let $\gamma_1 \cdots \gamma_5$ be the sequence of length $5 \cdot 5^k$ following l', which necessarily consists of α_k and β_k by Lem. 17.3. We prove that $\mathcal{A}$ must accept within this sequence. Let $q_0, q_1, \ldots, q_5$ with $q = q_0$ be the sequence of states such that $q_i \xrightarrow{\gamma_{i+1}} q_{i+1}$. Then $q_i \notin M$ for all $0 \leq i \leq 5$, for otherwise, by Lem. 17.3 and Lem. 18, $\mathcal{A}$ does not accept at all from l since the run gets trapped in M, which contradicts the assumption on acceptance from l. By Lem. 17.4 the sequence $\gamma_1 \cdots \gamma_5$ must contain two consecutive α_k followed by β_k or two consecutive β_k followed by α_k. W.l.o.g. suppose that two consecutive α_k are followed by β_k and that this concerns $\gamma_1, \gamma_2, \gamma_3$. If $q_0 \in A, q_1 \in A$ or $q_2 \in B$, we are done, since acceptance follows within the next γ_i. The remaining possibility is that all of $q_0, q_1 \in B \setminus A$ and $q_2 \in A \setminus B$ hold. However, this is not possible: Since $q_1 \xrightarrow{\alpha_k} q_2$ and $q_2 \in A$ implies that there must be $q' \in F$ such that $q_2 \xrightarrow{\alpha_k} q'$, we have that $q_1 \xrightarrow{\alpha_{k+1}} q'$ which implies $q_1 \in A$. Hence, $\mathcal{A}$ either accepts within $l' - l + 5 \cdot 5^k \leq 6 \cdot 5^k - 1$ steps from l, or does not accept from l at all. ◀

Putting this together with the results obtained in the previous section we obtain the following.

▶ **Corollary 20.** *w_∞ has finite convergence for $\mathcal{L}_\mu$ (but no finite bisimulation quotient).*

This follows from Cor. 16 and Lem. 17.7.

▶ Remark 21. Closer inspection of the proofs in this section yields two additional results. It follows from the proof of Lem. 18 that if $A_i = A_{i+1}$, then $A_i = A_j$ for all $j \geq i$. Hence, k can actually be computed effectively by computing the A_i and the B_i until both sequences stabilise. Moreover, the results of Lem. 18 and Thm. 19 do not rely on the exact form of w_∞, but rather its pattern, and the results from Sec. 4 do not make any assumptions on the word in question. It is possible to generalise the proof to words constructed via

$$\alpha'_0 = a \ , \quad \beta'_0 = b \ , \quad \alpha'_{i+1} = \alpha_i'^m \beta_i'^n \alpha_i'^m \quad , \quad \beta'_{i+1} = \beta_i'^m \alpha_i'^n \beta_i'^m$$

where $m > 1$. I.e. the importance is the symmetry between the α'_i and β'_i, as well as the use of two copies of α'_i at the beginning of α'_{i+1} etc. Moreover, this sequence of finite words does not have to be strictly monotonic in the use of its building blocks, i.e., the result also holds for words of the form $\alpha'_{i_0}\alpha'_{i_1}\cdots$, where $i_j \leq i_{j+1}$ for all $j \geq 0$. The case where the pattern eventually stabilises is not very interesting, of course, but bounded runs for TrA and, hence finite convergence of $\mathcal{L}_\mu$ still follow for the case where for all $k \in \mathbb{N}$ there is j such that $i_j \leq k$, i.e., the word uses α'_i of unbounded length as building blocks.

6 Infinite Convergence Through Higher-Order

We now show that the finite convergence of $\mathcal{L}_\mu$ formulas on the word w_∞ from Sec. 5 is due to the well-known fact that the expressive power of $\mathcal{L}_\mu$ is restricted to regular properties. In contrast, finite convergence does not hold anymore for a higher-order extension of $\mathcal{L}_\mu$ with non-regular expressiveness: Higher-Order Modal Fixpoint Logic (HFL). It extends $\mathcal{L}_\mu$ by the ability to form function definitions via λ abstraction. We refer to the literature or the appendix for a detailed introduction into HFL [22].

Let $\tau = (\mathbb{N} \to \mathbb{N}) \to (\mathbb{N} \to \mathbb{N}) \to \mathbb{N}$ be the set-theoretic type of functions that consume two functions of type $\mathbb{N} \to \mathbb{N}$ and return a natural number. Consider the HFL formula

$$\varphi = \big(\nu(X\colon \tau).\, \lambda(f, g\colon \mathbb{N} \to \mathbb{N}).\, f(\mathtt{tt}) \wedge X(f^2 \circ g \circ f^2, g^2 \circ f \circ g^2)\big)\,(\langle b\rangle, \langle a\rangle).$$

Here, $\psi = \lambda(f, g\colon \mathbb{N} \to \mathbb{N}).\, f(\mathtt{tt}) \wedge \ldots$ defines an anonymous function that takes as input two functions f and g of type $\mathbb{N} \to \mathbb{N}$ and returns the expression defined by the conjunction. The left conjunct, for example, defines the result of applying f to the set defined by $\mathtt{tt}$, i.e. $\mathbb{N}$. The formula $\langle a\rangle$ defines the function $S \mapsto \{i \mid w[i] = a \text{ and } i + 1 \in S\}$, and similarly for $\langle b\rangle$. The fixpoint X itself is now of higher-order, and it is equivalent to the expression $\bigwedge_{i\in\mathbb{N}} \psi^i$, where $\psi^0 = \lambda(f, g\colon \mathbb{N} \to \mathbb{N}).\, \mathtt{tt}$ and $\psi^{i+1} = \psi[\psi^i/X]$. Applying this expression to the original arguments $\langle b\rangle, \langle a\rangle$ in φ and using some β-reduction, we obtain that φ is equivalent to

$$\psi^0(\langle b\rangle, \langle a\rangle) \wedge \psi^1(\langle b\rangle, \langle a\rangle) \wedge \psi^2(\langle b\rangle, \langle a\rangle) \wedge \ldots.$$

With standard arguments about λ-expressions and modal logic, we obtain that $\psi^i(\langle b\rangle, \langle a\rangle)$ defines the set of positions such that all the α_j for $j < i$ follow, where α_j is as in Sec. 5. Hence, we can conclude that φ defines the set of positions i in w_∞ of Sec. 5 such that the postfix following i starts with α_j for all $j \in \mathbb{N}$.

▶ **Theorem 22.** *The extension* HFL^2 *of* $\mathcal{L}_\mu$ *by second-order functions does not have finite convergence on* w_∞*.*

Proof. Following the argument above, we get $[\![\varphi]\!]^{w_\infty} = \emptyset$. However, from the previous analysis we can see that a position $\sum_{i=0}^{j-1} 5^i$, namely the starting point of the first α_j in w_∞, is still contained in the jth approximant $(\bigwedge_{0\le i\le j} \psi^i(\langle b\rangle, \langle a\rangle)$ but not in the $j+1$st one. Hence, φ does not have finite convergence on w_∞. ◀

▶ Remark 23. It is possible to strengthen Thm. 22 by constructing a formula in HFL^1, the first-order extension of $\mathcal{L}_\mu$. (We have $\mathcal{L}_\mu = \mathrm{HFL}^0$.) But the construction is complicated and requires further insight into the semantics of HFL, so it is left out for space considerations.

7 Conclusion

We have presented a further contribution to the theory of closure ordinals for μ-calculi, namely the – possibly surprising – fact that having finite bisimulation quotients is not an equivalent but a strictly stronger property than having finite fixpoint convergence, and that this discrepancy is due to $\mathcal{L}_\mu$'s relatively restricted expressive power. As a corollary, we obtain decidability of $\mathcal{L}_\mu$ model checking over w_∞ (and the class of words built like it), i.e. for given φ and i, it is decidable whether φ holds at position i in w_∞. This follows since every $\mathcal{L}_\mu$ formula is equivalent to one in ML over w_∞ and it is easily decidable whether the letter at some position is an a or a b. The use of the Compactness Theorem in Lem. 11 might look prohibitively non-constructive, but the result follows from the constructive nature of the proof up to alternation-free $\mathcal{L}_\mu$ (cf. Rem. 12). Note that this result does not follow from results around morphic words (cf. e.g. [2]), since w_∞ is not morphic. Hence, the construction of w_∞ can be used as the basis for a new class of infinite structures with decidable $\mathcal{L}_\mu$ model checking beyond pushdown processes. However, the design follows a pattern quite similar to morphic words, and we plan to investigate possible links between the two concepts.

There are some other directions into which our research can be extended: for further technical developments the theory of finite convergence has been formulated over classes of structures, even though it has been used here for a single word structure only. It remains to be seen how far the construction pattern can be stretched, i.e. what a largest class of structures with finite convergence is. One may leave the world of words without losing the ability to reduce from the entire $\mathcal{L}_\mu$ to the alternation-free fragment as there are richer classes of structures with corresponding collapse of the alternation hierarchy [14].

The type hierarchy in HFL – which is strict in terms of expressiveness [3] – gives rise to the question whether for each level i, there is a class of structures on which HFL^i has finite convergence but HFL^{i+1} does not. Here, we have answered the question for $i = 0$. For $i > 0$ this is tricky as these higher-order logics lack comparable automata-theoretic support, cf. [6].

References

1 B. Afshari and G. 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*, volume 23 of *LIPIcs*, pages 30–44. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2013. doi:10.4230/LIPIcs.CSL.2013.30.

2 J.-P. Allouche and J. O. Shallit. *Automatic Sequences - Theory, Applications, Generalizations*. Cambridge University Press, 2003.

3 R. Axelsson, M. Lange, and R. Somla. The complexity of model checking higher-order fixpoint logic. *Logical Methods in Computer Science*, 3:1–33, 2007. doi:10.2168/LMCS-3(2:7)2007.

4 J. Barwise and Y. N. Moschovakis. Global inductive definability. *J. Symb. Log.*, 43(3):521–534, 1978. doi:10.2307/2273529.

5 J. C. Bradfield. The modal μ-calculus alternation hierarchy is strict. In *Proc. 7th Conf. on Concurrency Theory, CONCUR'96*, volume 1119 of *LNCS*, pages 233–246. Springer, 1996. doi:10.1007/3-540-61604-7_58.

6 F. Bruse. *Extremal Fixpoints for Higher-Order Modal Logic*. PhD thesis, University of Kassel, Germany, 2020. doi:doi:10.17170/kobra-202008091547.

7 F. Bruse, O. Friedmann, and M. Lange. On guarded transformation in the modal μ-calculus. *Log. J. IGPL*, 23(2):194–216, 2015. doi:10.1093/jigpal/jzu030.

8 F. Bruse, M. Lange, and É. Lozes. Collapses of fixpoint alternation hierarchies in low type-levels of higher-order fixpoint logic. Proc. Workshop on Programming and Reasoning on Infinite Structures, PARIS'14, 2018.

9 M. 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*, pages 35–39. Laboratoire d'Informatique Fondamentale de Marseille, 2010. URL: https://hal.archives-ouvertes.fr/hal-00512377/document#page=36.

10 Mads Dam. Ctl* and ectl* as fragments of the modal mu-calculus. *Theor. Comput. Sci.*, 126(1):77–96, 1994. doi:10.1016/0304-3975(94)90269-0.

11 S. Ghilardi, M. J. Gouveia, and L. Santocanale. Fixed-point elimination in the intuitionistic propositional calculus. *ACM Trans. Comput. Log.*, 21(1):4:1–4:37, 2020. doi:10.1145/3359669.

12 M. J. Gouveia and L. Santocanale. $\aleph_1$ and the modal μ-calculus. *Log. Methods Comput. Sci.*, 15(4), 2019. doi:10.23638/LMCS-15(4:1)2019.

13 D. Janin and I. Walukiewicz. On the expressive completeness of the propositional μ-calculus with respect to monadic second order logic. In *Proc. 7th Conf. on Concurrency Theory, CONCUR'96*, volume 1119 of *LNCS*, pages 263–277. Springer, 1996. doi:10.1007/3-540-61604-7_60.

14 J.Gutierrez, F.Klaedtke, and M. Lange. The μ-calculus alternation hierarchy collapses over structures with restricted connectivity. *Theor. Comput. Sci.*, 560:292–306, 2014. doi:10.1016/j.tcs.2014.03.027.

15 R. Kaivola. Axiomatising linear time mu-calculus. In *CONCUR '95: Concurrency Theory, 6th International Conference, Philadelphia, PA, USA, August 21-24, 1995, Proceedings*, volume 962 of *Lecture Notes in Computer Science*, pages 423–437. Springer, 1995. doi:10.1007/3-540-60218-6_32.

16 S. C. Kleene. On a notation for ordinal numbers. *Journal of Symbolic Logic*, 3:150–155, 1938.
17 D. Kozen. Results on the propositional μ-calculus. *TCS*, 27:333–354, 1983. doi:10.1016/0304-3975(82)90125-6.
18 E. Lozes. A type-directed negation elimination. In *Proceedings Tenth International Workshop on Fixed Points in Computer Science, FICS 2015, Berlin, Germany, September 11-12, 2015*, volume 191 of *EPTCS*, pages 132–142, 2015. doi:10.4204/EPTCS.191.12.
19 G. C. Milanese and Y. 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*, volume 11541 of *Lecture Notes in Computer Science*, pages 498–515. Springer, 2019. doi:10.1007/978-3-662-59533-6_30.
20 M. 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*, volume 1563 of *Lecture Notes in Computer Science*, pages 531–540. Springer, 1999. doi:10.1007/3-540-49116-3_50.
21 R. S. Streett and E. A. Emerson. The propositional mu-calculus is elementary. In *Automata, Languages and Programming, 11th Colloquium, Antwerp, Belgium, July 16-20, 1984, Proceedings*, volume 172 of *Lecture Notes in Computer Science*, pages 465–472. Springer, 1984. doi:10.1007/3-540-13345-3_43.
22 M. Viswanathan and R. Viswanathan. A higher order modal fixed point logic. In *Proc. 15th Int. Conf. on Concurrency Theory, CONCUR'04*, volume 3170 of *LNCS*, pages 512–528. Springer, 2004. doi:10.1007/978-3-540-28644-8_33.
23 T. Wilke. Alternating tree automata, parity games, and modal μ-calculus. *Bull. Belgian Math. Soc.*, 8(2):359–391, 2001. doi:10.36045/bbms/1102714178.

A Proof of Lemma 7

▶ **Lemma 7.** *Let φ be an $\mathcal{L}_\mu$ formula. Then the bottom-up unfolding of φ (cf. Def. 2) and the top-down unfolding of φ (cf. Def. 5) are equivalent to the unfolding defined in Def. 6. In particular, the top-down unfolding is well-defined.*

Proof. We first show that the top-down unfolding of φ is equivalent to the unfolding defined in Def. 6. W.l.o.g. let $\varphi = \sigma_{j+1}X_{j+1}.\psi_{j+1}$ be a formula with fixpoint variables $X_1, \dots, X_k$ such that at most $X_1, \dots, X_j$ for $j \leq k-1$ appear freely in φ. Let $n \geq 0$. We show that

$$(\psi_{j+1}[\psi_1/X_1, \dots, \psi_j/X_j])^{(i,n,\dots,n)} \equiv (\psi^i_{j+1}[\psi_1/X_1, \dots, \psi_j/X_j])^n \tag{1}$$

for closed and fixpoint-free formulas $\psi_1, \dots, \psi_j$ and $i \leq n$. Note that the left term refers to the unfolding according to Def. 6, while the right one refers to the ith unfolding of the single fixpoint X_{j+1}, followed by the nth top-down unfolding of all remaining fixpoint-subformulas. The claim of the Lemma regarding the top-down unfolding procedure then follows with $j = 0$ and $i = n$.

We start with an induction over j. For the base case $j = k-1$ it follows that ψ_j does contain only one fixpoint-subformula and together with the observation that the $\psi_1, \dots, \psi_j$ are closed and fixpoint-free formulas, both sides of Eq. 1 coincide with a single fixpoint unfolding. Let $j < k-1$ and assume that the result has been shown for all j' with $j < j' \leq k-1$ and for all $i \leq n$. We show the result for j by induction on i. For $i = 0$, we have that $\psi^0_{j+1} = \hat{\sigma}_{j+1}$ and, hence $(\psi^0_{j+1}[\psi_1/X_1, \dots, \psi_j/X_j])^n = \hat{\sigma}_{j+1}$ and $(\psi_{j+1}[\psi_1/X_1, \dots, \psi_j/X_j])^{(0,n,\dots,n)} = \hat{\sigma}_{j+1}$ as well. Next let $i > 0$ and assume that the result has been proven for all $i' \leq i$. We show this case by another induction on the structure of ψ_{j+1}. The base case a, as well as the boolean cases and $\bigcirc$ are straightforward. The interesting cases are the base case of a fixpoint variable $X_{j'}$ for $j' \leq j$, the case of X_{j+1} itself and the case of a fixpoint definition of the form $\sigma_{j'} X_{j'}.\psi_{j'}$ for $j' > j+1$. The first case results in the same subformula, because in both versions of the unfolding $\psi_{j'}$ is substituted for $X_{j'}$ before the unfolding, and this is a closed and fixpoint-free formula. For X_{j+1} itself, we use the induction hypothesis of the induction over i, namely that $(X_{j+1}[\psi_1/X_1, \ldots, \psi_j/X_j])^{(i,n,\ldots,n)} = (\psi_{j+1}[\psi_1/X_1, \ldots, \psi_j/X_j])^{(i-1,n,\ldots,n)}$ is equivalent to $(\psi_{j+1}^{i-1}[\psi_1/X_1, \ldots, \psi_j/X_j])^n$. For the case of a formula of the form $\sigma_{j'} X_{j'}.\psi_{j'}$, we use the induction hypothesis for the induction over j, namely that

$$(\psi_{j'}[\psi_1/X_1, \ldots, \psi_j/X_j, \psi'_{j+1}/X_{j+1}, \psi'_{j+2}/X_{j+2}, \ldots, \psi'_{j'-1}/X_{j'-1}])^{(n,\ldots,n)} \equiv (\psi_{j'}^n[\psi_1/X_1, \ldots, \psi_j/X_j, \psi''_{j+1}/X_{j+1}, \psi''_{j+2}/X_{j+2}, \ldots, \psi''_{j'-1}/X_{j'-1}])^n$$

for $\psi'_{j+1} = (\psi_{j+1}[\psi_1/X_1, \ldots, \psi_j/X_j])^{(i-1,n,\ldots,n)}$ and $\psi''_{j+1} = (\psi_{j+1}^{i-1}[\psi_1/X_1, \ldots, \psi_j/X_j])^n$, which, by the induction hypothesis of the induction over i, are equivalent.[1] This finishes the induction over ψ_{j+1} and with it, the induction over i, and the induction over j.

For the bottom-up unfolding of Def. 2, we show equivalence to the unfolding from Def. 6 by showing that unfolding an innermost formula will not change the formula generated by either procedure. The result then follows by an induction over the ordering of fixpoints used in the bottom-up unfolding. Let φ be a fixpoint formula, and let $X_1, \ldots, X_k$ be an enumeration of its fixpoint variables such that $X_i \not<_{\text{fp}} X_j$ for $j > i$, and let $\sigma_i X_i.\,\psi_i$ be the defining formula of X_i. Let $\varphi[\psi_k^n/\sigma_k X_k.\,\psi_k]$ be the formula obtained by unfolding X_k n times. If we can show that, for all subformulas ψ of φ, and for all $s = (i_1, \ldots, i_{k-1}, n)$, we have $\psi^s \equiv (\psi[\psi_k^n/\sigma_k X_k.\,\psi_k])^s$, we are done. We show this by induction on the lexicographical ordering of the tuple, i.e starting with $s = (0, \ldots, 0, n)$, for which the result clearly holds. Assume that we have shown it for all s' that are lexicographically smaller than s. We show the result by another induction on ψ. The base case a, the boolean cases, and the case $\bigcirc$ are straightforward. The interesting cases are a variable X_j for $j < k$, a fixpoint definition of the form $\sigma_j X_j.\psi_j$ for $j < k$, and the fixpoint definition of the form $\sigma_k X_k.\psi_k$. For the case of X_j, if $s(j) = 0$, the result is immediate. For the case that $s(j) > 0$, and for the case of $\sigma_j X_j.\psi_j$, the formula to be substituted is $\text{fp}(X_j)^{s[j--]}$, respectively $(\text{fp}(X_j)[\psi_k^n/\sigma_k X_k.\psi_k])^{s[j--]}$, for which the result holds since $s[j--]$ is lexicographically smaller than s. For the case $\sigma_k X_k.\psi_k$, there are two possibilities. If $s(j) = 0$ for all $0 \le j < k$, then the result follows from the base case. If not, note that X_k is minimal w.r.t. $<_f$ p and, hence, its defining formula $\text{fp}(X_k) = \psi_k$ contains no fixpoint definitions itself. However, it can contain free fixpoint variables from among the $X_1, \ldots, X_{k-1}$. We first note that, by definition, $\text{fp}(X_k)^s$ for $s = (i_1, \ldots, i_{k-1}, n)$ is equivalent to $\psi_k^n[\psi_1/X_1, \ldots, \psi_{k-1}/X_{k-1}]$, where ψ_j is $\hat{\sigma}_j$ if $s[j] = 0$ and $\text{fp}(X_j)^{s[j--]}$ otherwise. But, using the induction hypothesis for the ψ_i, this is exactly $(\sigma_k X_k.\psi_k[\psi_k^n/\sigma_k X_k.\psi_k])[\psi_1/X_1, \ldots, \psi_{k-1}/X_{k-1}]$. This finishes the proof. ◀

B Proof of Lemma 13

▶ **Lemma 13.** *Let w be an infinite word. If w has finite convergence for the set of all unipolar $\mathcal{L}_\mu$ formulas, it has finite convergence for the set of all alternation-free $\mathcal{L}_\mu$ formulas.*

Proof. Let φ an alternation-free $\mathcal{L}_\mu$ formula. The set of fixpoint variables of φ can be partitioned into the sets $\mathcal{X}_1, \ldots, \mathcal{X}_k$ as described in Sec. 2. First, we show by induction on these sets that all fixpoint-subformulas of φ can be unfolded starting from $\mathcal{X}_k$ and ending with $\mathcal{X}_1$ while preserving equivalence over w. First, note that for all $X \in \mathcal{X}_k$ there is no fixpoint-subformula in $\text{fp}(X)$ that has a different polarity. This implies that all fixpoint-subformulas of

[1] The equality $\psi'_l = \psi''_l$ for $j+2 \le l \le j'-1$ can also be inferred via induction for a previous j'.

φ with $X \in \mathcal{X}_k$ are unipolar, but not necessarily closed. If we take a fixpoint-subformula with $X \in \mathcal{X}_k$ such that there is no $X' \in \mathcal{X}_k$ with $X' <_{\mathrm{fp}} X$ it follows from the alternation-freeness of φ that $\mathrm{fp}(X)$ is also closed and, thus, by the assumption of the lemma that there is an equivalent unfolding over w. If we replace all such fixpoint-subformulas with their respective equivalent unfolding we have unfolded all fixpoint-subformulas with $X \in \mathcal{X}_k$. Under the assumption that all fixpoint subformulas with $X \in \mathcal{X}_i$ are already unfolded, we can infer with the same arguments that there is an equivalent unfolding for all fixpoint subformulas of $\mathcal{X}_{i-1}$. By the principle of induction this shows that there is some equivalent finite unfolding of φ over w. What is left to argue is that there is a uniform one, i.e., that there is n such that φ^n is equivalent to φ. Note that for each closed, unipolar fixpoint subformula $\sigma X.\psi$ with equivalent unfolding ψ^m it holds that $\psi^{m'}$ with $m' \geq m$ is equivalent as well. As the number of fixpoint subformulas in φ is finite, n is given by the maximum number of unfoldings needed for any fixpoint subformula. ◀

C Additional Material for Section 6

We give a brief introduction to HFL on words. A more thorough introduction for HFL on arbitrary LTS can be found in e.g. [22]. We also point out that the exposition in Sect. 6 makes use of syntactic sugar that is not directly covered by the pure syntax. The constructs, however, are all straightforward using only principles which are standard in λ-calculi. For instance further below we explain that the subformula $\langle a \rangle$ is used to abbreviate something like $\lambda x.\langle a \rangle x$, so this simply makes use of η-conversion.

Types. Consider the set of types defined via

$$\tau ::= \bullet \mid \tau \to \tau.$$

The type $\bullet$ is the base type of subsets of $\mathbb{N}$, for example those defined by an $\mathcal{L}_\mu$ formula. A type $\tau_1 \to \tau_2$ is inhabited by monotone functions from τ_1 to τ_2.

Such a type induces a lattice over a given word w via the following definition

$$\llbracket \bullet \rrbracket^w = (\mathcal{P}(\mathbb{N}), \subseteq)$$
$$\llbracket \tau_1 \to \tau_2 \rrbracket^w = (\llbracket \tau_2 \rrbracket^w \to \llbracket \tau_1 \rrbracket^w, \sqsubseteq_{\tau_1 \to \tau_2})$$

where $\llbracket \tau_2 \rrbracket^w \to \llbracket \tau_1 \rrbracket^w$ denotes the set of functions from $\llbracket \tau_1 \rrbracket^w$ to $\llbracket \tau_2 \rrbracket^w$ and $\sqsubseteq_{\tau_1 \to \tau_2}$ is defined as the pointwise order via $f \sqsubseteq_{\tau_1 \to \tau_2} g$ iff $f(x) \sqsubseteq_{\tau_2} g(x)$ for all $x \in \llbracket \tau_1 \rrbracket^w$. Here, $\sqsubseteq_\bullet$ ordinary set inclusion $\subseteq$. All these lattices are complete since $\llbracket \bullet \rrbracket^w$ is complete, and a lattice of functions is complete if the functions are into a complete lattice and ordered pointwise.

Syntax. Let $\mathcal{V} = \{x, y, \dots\}$ be a set of λ variables. The syntax of HFL extends that of $\mathcal{L}_\mu$ to

$$\varphi ::= a \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \bigcirc \varphi \mid X \mid x \mid \mu(X\colon \tau).\varphi \mid \nu(X\colon \tau).\varphi \mid \lambda(x\colon \tau).\varphi \mid \varphi\, \varphi$$

where τ is a type.

The intuition behind the operators that are new in comparison to $\mathcal{L}_\mu$ is as follows.

- $\lambda x.\varphi$ defines an anonymous function that consumes an argument, bound to x, and returns the value of φ if x is set to that argument.
- $\varphi\, \psi$ denotes the application of ψ to φ, and
- the fixpoints can now be of higher order, i.e. define function type objects.

$$\frac{}{\Gamma \vdash a\colon \bullet} \qquad \frac{}{\Gamma, X\colon \tau \vdash X\colon \tau} \qquad \frac{}{\Gamma, x\colon \tau \vdash x\colon \tau} \qquad \frac{\Gamma \vdash \varphi_1\colon \bullet \qquad \Gamma \vdash \varphi_2\colon \bullet}{\Gamma \vdash \varphi_1 \vee \varphi_2\colon \bullet}$$

$$\frac{\Gamma \vdash \varphi_1\colon \bullet \qquad \Gamma \vdash \varphi_2\colon \bullet}{\Gamma \vdash \varphi_1 \wedge \varphi_2\colon \bullet} \qquad \frac{\Gamma \vdash \varphi\colon \bullet}{\Gamma \vdash \bigcirc \varphi\colon \bullet} \qquad \frac{\Gamma, X\colon \tau \vdash \varphi\colon \tau'}{\Gamma \vdash \lambda(X\colon \tau).\, \varphi\colon \tau \to \tau'} \qquad \frac{\Gamma, X\colon \tau \vdash \varphi\colon \tau}{\Gamma \vdash \mu(X\colon \tau).\, \varphi\colon \tau}$$

$$\frac{\Gamma, X\colon \tau \vdash \varphi\colon \tau}{\Gamma \vdash \nu(X\colon \tau).\, \varphi\colon \tau} \qquad \frac{\Gamma \vdash \varphi\colon \tau \to \tau' \qquad \Gamma \vdash \varphi'\colon \tau}{\Gamma \vdash \varphi\, \varphi'\colon \tau'}$$

Figure 1 The type system of HFL.

Note that we have chosen not to introduce negation. HFL admits negation normal form [18], whence negation only needs to occur on front of atomic propositions. Over words, the formula $\neg a$ however is equivalent to $\bigvee_{b\in\Sigma, b\neq a} b$ as stated for $\mathcal{L}_\mu$ already, so negation can be avoided altogether.

An advantage of this avoidance of negation is the simplification of the type system to monotone functions only. The original definition of HFL includes negation as a syntactic construct at the expense of a slightly more complex type systems which needs to keep track of antitonicity information so that fixpoints are guaranteed to exists due to monotonicity.

Without this, the only purpose of the type system is to avoid misapplications as in $a\, b$ for instance which cannot be given proper meaning. Another example is $a \vee \lambda(x\colon \mathbb{N}).\, x \wedge b$.

In the absence of negation, an HFL formula φ is said to be well-typed if the statement $\emptyset \vdash \varphi$ can be derived via the rules shown in Fig. 1. The sequence Γ on the left of a typing statement is called a typing context and collects typing hypotheses.

Semantics. In order to endow well-typed HFL formulas with semantics, we extend environments to $\mathcal{V}$, i.e. environments can also store values for λ-variables which may be higher-order objects depending on the type of the variable. The semantics of an HFL formula φ is given inductively as per the following:

$$\begin{aligned}
[\![a]\!]^w_\eta &= \{i \in \mathbb{N} \mid w[i] = a\}\\
[\![X]\!]^w_\eta &= \eta(X)\\
[\![x]\!]^w_\eta &= \eta(x)\\
[\![\varphi \vee \psi]\!]^w_\eta &= [\![\varphi]\!]^w_\eta \cup [\![\psi]\!]^w_\eta\\
[\![\varphi \wedge \psi]\!]^w_\eta &= [\![\varphi]\!]^w_\eta \cap [\![\psi]\!]^w_\eta\\
[\![\bigcirc \varphi]\!]^w_\eta &= \{i \in \mathbb{N} \mid i+1 \in [\![\varphi]\!]^w_\eta\}\\
[\![\lambda(x\colon \tau).\varphi]\!]^w_\eta &= f \in [\![\tau \to \tau']\!]^w \quad \text{where f.a. } d \in [\![\tau]\!]^w.f(d) = [\![\varphi]\!]^w_{\eta[x\mapsto d]}\\
&\ \text{with } \tau' \text{ the type of } \varphi\\
[\![\varphi\, \psi]\!]^w_\eta &= [\![\varphi]\!]^w_\eta([\![\psi]\!]^w_\eta)\\
[\![\mu(X\colon \tau).\varphi]\!]^w_\eta &= \bigsqcap\{d \subseteq [\![\tau]\!]^w \mid [\![\varphi]\!]^w_{\eta[X\mapsto d]} \sqsubseteq_\tau d\}\\
[\![\nu(X\colon \tau).\varphi]\!]^w_\eta &= \bigsqcup\{d \subseteq [\![\tau]\!]^w \mid d \sqsubseteq_\tau [\![\varphi]\!]^w_{\eta[X\mapsto d]}\}
\end{aligned}$$

The formula φ from Sect. 6. We give some additional explanation about the formula φ defined in Sec. 6. Recall that

$$\varphi = \big(\nu(X\colon \tau).\, \lambda(f, g\colon \mathbb{N} \to \mathbb{N}).\, f(\texttt{tt}) \wedge X(f^2 \circ g \circ f^2, g^2 \circ f \circ g^2)\big)\,(\langle b\rangle, \langle a\rangle).$$

Notation such as $\lambda(f, g\colon \mathbb{N} \to \mathbb{N}).\ \ldots$ can easily be seen to be an abbreviation of the longer $\lambda(f\colon\ \ldots).\,\lambda(g\colon\ \ldots).$, and the same goes for the application to X written in a similar style. Clearly, neither function composition nor the function $\langle a\rangle$ are in the official syntax of HFL. The function $\langle a\rangle$ can be written in standard syntax as $\lambda(x\colon \mathbb{N}).\, a \wedge \bigcirc x$. Function composition, here between five functions, is an abbreviation for $\lambda(x\colon \mathbb{N}).\, f(f(g(f(f\,x))))$, respectively $\lambda(x\colon \mathbb{N}).\, g(g(f(g(g\,x))))$.

The intuition given for the semantics of φ is already close to the true semantics of φ. Since the Kleene Fixpoint Theorem applies in this setting, too, we can use it to obtain the semantics of φ on w_∞. As in Sec. 6, write ψ for the defining formula of $\lambda f, g.\, f(\texttt{tt}) \wedge X(f^2 \circ g \circ f^2, g^2 \circ f \circ g^2)$, write ψ^0 for $\lambda f, g.\, \texttt{tt}$ and ψ^{i+1} for $\psi[\psi^i/X]$. Given some functions f and g, we introduce the following functions:

$$f_1(x) = f(x) \qquad f_{i+1} = f_i(f_i(g_i(g_i(g_i(x)))))$$
$$g_1(x) = g(x) \qquad g_{i+1} = g_i(g_i(f_i(g_i(g_i(x)))))$$

Then ψ^i defines the function $f, g \mapsto f_1(\texttt{tt}) \wedge f_2(\texttt{tt}) \wedge \cdots \wedge f_i(\texttt{tt})$ by induction on i. Hence, the semantics of X is $f, g \mapsto \bigwedge_{i\in\mathbb{N}} f_i(\texttt{tt})$. Applied to the arguments $\langle a\rangle$ and $\langle b\rangle$, this yields, by abuse of notation, $\bigwedge_{i\in\mathbb{N}}\langle \alpha_i\rangle\texttt{tt}$, which is as claimed in Sec 6. The remarks on the infinite convergence process of φ on w_∞ then follow.