Collapses of Fixpoint Alternation Hierarchies in Low Type-Levels of Higher-Order Fixpoint Logic

Florian Bruse[1], Martin Lange[1], and Etienne Lozes[2]

[1] School of Electr. Eng. and Computer Science, University of Kassel, Germany
[2] I3S, Université Nice Sophia Antipolis, France

1 Higher-Order Fixpoint Logic

Higher-Order Fixpoint Logic (HFL) [10] is an extension of the modal μ-calculus $\mathcal{L}_\mu$ [7] by means of a typed λ-calculus. Its formulas not only denote predicates but also higher-order objects like predicate transformers etc. Semantics is given via higher-order functions, ordered pointwise to form complete lattices so that least and greatest fixpoint quantifiers can also be used over higher-order predicates.

HFL obtains great expressive power in this way, as also indicated by its complexities: model checking is non-elementary [1], satisfiability is undecidable [9, 10]. Its lack of the finite model property (FMP) makes it inherently a logic for infinite-state systems; there is a close correspondence between HFL model checking on finite-state systems and $\mathcal{L}_\mu$ model checking on recursion schemes [6]. This uses principles which, intuitively, swap non-regularity on the model side for non-regularity on the logic side. This cannot be done straight-forwardly because of the much richer fixpoint quantifier structure in HFL formulas compared to the plain greatest fixpoint semantics in recursion schemes. In order to make the link between higher-order model checking and model checking higher-order fixpoint logic even tighter we aim to understand the nature of fixpoint alternation in HFL, i.e. the effects on expressiveness for formulas with entangled fixpoint quantifiers of different kinds.

It turns out that the picture for HFL is a lot more complex than that for $\mathcal{L}_\mu$ where the alternation hierarchy is known to be strict in general [2] and to collapse only over classes of "simple" structures, c.f. [5, 4]. The type-order of functions plays a dominant role in questions after the hierarchy. It is strict up to HFL^2, where HFL^k denotes the fragment that uses functions of order at most k [3]. Due to bisimulation-invariance, strictness already holds over the class of trees but not necessarily over finite models due to the lack of the FMP.

This paper presents some principles which can be used to eliminate fixpoint alternation at the expense of an increase in type order. They are, however, only applicable over particular classes of (infinite) structures. Let $\mathbb{T}^{\sim}_{\mathsf{fin}}$ denote the class of transition systems that have finite bisimulation quotients and, for all $k \geq 0$, $\mathbb{T}^{k}_{\mathsf{fin}}$ the class of transition systems over which all fixpoint iterations definable in HFL^i stabilise after finitely many steps. It is not hard to see that we have

$$\mathbb{T}^{0}_{\mathsf{fin}} \supseteq \mathbb{T}^{1}_{\mathsf{fin}} \supseteq \cdots \supseteq \bigcap_{i \in \mathbb{N}} \mathbb{T}^{i}_{\mathsf{fin}} \supseteq \mathbb{T}^{\sim}_{\mathsf{fin}} . \tag{1}$$

We briefly sketch the logic HFL. For a more detailed description, including a formal definition of the semantics via pointwise ordered monotone functions in complete lattices, see [10]. HFL incorporates a simple type system with a single base type $\bullet$, interpreted via sets of states in a transition system, and a binary type constructor $\to$ for building functions. The *order* of a type is defined via $ord(\bullet) = 0$ and $ord(\sigma \to \tau) = \max\{ord(\tau), 1 + ord(\sigma)\}$. HFL admits negation normal form (NNF) [8]. We use this here and introduce the syntax without an explicit negation operator, but we shall be cavalier about this for the sake of readability. *Formulas* are of the form

$$\varphi \;:=\; p \mid \overline{p} \mid X \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \langle a \rangle \varphi \mid [a]\varphi \mid \mu X^{\tau}.\varphi \mid \nu X^{\tau}.\varphi \mid \lambda X^{\tau}.\varphi \mid \varphi\;\varphi$$

where X is a variable, p is an atomic proposition interpreted by a set of states in an LTS, a is an action interpreted as a set of edges in an LTS, and τ is a type. The fragment HFL^k, $k \geq 0$, consists of all formulas of type $\bullet$ which use types of order at most k.

The formula $\lambda f^{\bullet \to \bullet}.\lambda X^{\bullet}.f\ \langle a \rangle X$ for instance has type $(\bullet \to \bullet) \to \bullet \to \bullet$ and is therefore not of the base type itself. Hence, it can only occur as a subformula where it can be used to modify a predicate transformer f by prepending a modal diamond-operation to its arguments. The formula $(\mu f^{\bullet \to \bullet}.\lambda X^{\bullet}.X \vee (f\ [a]X) \vee (f\ [b]X))\ \overline{p}$ belongs to HFL^1. Using the principles of fixpoint unfolding and β-reduction one can see that it is equivalent to the formula $\bigvee_{w \in \{a,b\}^*}[w]\overline{p}$ of infinitary modal logic (which describes the non-universality problem for NFA).

For a last example, suppose $\varphi(X)$ is a $\mathcal{L}_\mu$- or, equivalently HFL^0-formula with a free predicate variable X. Note that its occurrences are necessarily positive if φ is in NNF. Consider the HFL^1 formula[1]

$$\mathsf{GFP}^0_\varphi := \big(\mu f^{\bullet \to \bullet}.\lambda X^{\bullet}.(X \wedge \Box^*(X \to \varphi(X))) \vee f\ \varphi(X)\big)\ \top$$

where $\Box^*\psi := \nu Y^{\bullet}.\psi \wedge \bigwedge_a [a]Y$ checks that ψ holds on all reachable states. Again, using fixpoint unfolding, one can see that GFP^0_φ checks whether there is some $n \in \mathbb{N}$ such that the current state s satisfies $\varphi^n(\top)$ and all reachable states satisfy $\varphi^n(\top) \to \varphi^{n+1}(\top)$, indicating that the nth approximation of the fixpoint coincides with the $(n+1)$th on those states that are reachable from the current one. By bisimulation-invariance, non-reachable states do not contribute to the fixpoint's value. Hence, this shows that greatest fixpoints at type order 0 which can be approximated with finitely many iterations, can be replaced *uniformly* by least fixpoints at type order 1 (with an embedded but not entangled greatest fixpoint formula at type order 0).

Theorem 1. *The* HFL^0*- or* $\mathcal{L}_\mu$*-alternation hierarchy collapses to the alternation-free fragment in* HFL^1 *over* $\mathbb{T}^0_{\mathsf{fin}}$.

2 The Collapse of the Order-1 Hierarchy in HFL^2

Let $\lambda x^{\bullet}.\varphi(F)$ be an HFL^1-formula with a free predicate F of type $\bullet \to \bullet$. Consider the formula

$$\mathsf{GFP}^1_\varphi := \Big(\mu g^{(\bullet \to \bullet) \to \bullet \to \bullet}.\lambda F^{\bullet \to \bullet}\lambda x^{\bullet}.\big(F\,x \wedge H\,(\lambda z^{\bullet}.F\,z \to (\lambda x^{\bullet}.\varphi(F))\,z)\big) \vee (g\,\varphi(F))\,x\Big)\ \lambda y^{\bullet}.\top$$

where $H^{(\bullet \to \bullet) \to \bullet}$ will be defined below. It does not contain $\lambda x^{\bullet}.\varphi(F)$ or g. Hence, GFP^1_φ is alternation-free if φ is so. It follows the pattern of GFP^0_φ encoding Kleene's Theorem: to compute the greatest fixpoint of the order-2 function $F \mapsto \lambda x^{\bullet}.\varphi(F)$, start with the top element of the respective lattice, namely $\lambda y^{\bullet}.\top$, and apply $\lambda x^{\bullet}.\varphi(F)$ until the iteration stabilises. This is what the right disjunct $(g\,\varphi(F))\,x$ of the formula above does. The conjunction on the left returns the value of the fixpoint at the argument x via $F\,x$ and verifies that the fixpoint iteration has actually stabilised. This is more complex now; in order to check that a $\mathcal{L}_\mu$ fixpoint has stabilised it suffices to encode set inclusion. An HFL^1 fixpoint is a predicate transformer, though. To check that its iteration has stabilised, we need to verify that two consecutive approximations agree *as functions*, i.e. on all possible arguments, which are all subsets of the underlying LTS. Enumerating all subsets of an LTS is impossible, though, since some may not be HFL-definable.

This problem can be resolved in the following way.

[1]The implicit negation in the use of $\to$ can be eliminated, but at significant cost to readability.

Lemma 2. *Let $\mathcal{T} \in \mathbb{T}^1_{\mathsf{fin}}$, $\nu F.\lambda X.\varphi$ be an* HFL^1*-fixpoint formula of type $\bullet \to \bullet$, and $F^i_{\mathcal{T}}$ denote its i-th approximant over $\mathcal{T}$ for $i \in \mathbb{N}$. If $F^{i+1}_{\mathcal{T}}(S) \supseteq F^i_{\mathcal{T}}(S)$ for all state sets S definable in modal logic, then the inclusion also holds for arbitrary* HFL^1*-definable state sets.*

Intuitively, this holds because the value of a greatest fixpoint on some HFL^k-definable argument depends on the value of the fixpoint on some other arguments which are accessed via HFL^k-formulas. Hence, they must be HFL^k-definable, too. Thus, it is sufficient to test for stabilisation only on HFL^k-definable arguments. For $k \geq 2$ these are not easy to enumerate either because of the use of fixpoint formulas and unbounded variables. Over each fixed structure in $\mathbb{T}^1_{\mathsf{fin}}$, though, any HFL^1 formula is equivalent to a fixpoint-free formula since all fixpoint definitions can be replaced by some finite approximation. Using β-reduction, we even get equivalence to a formula of modal logic (ML).

Now consider the fixpoint formula $H^{(\bullet\to\bullet)\to\bullet}$ defined as

$$H = \nu H^{(\bullet\to\bullet)\to\bullet}.\, \lambda t^{\bullet\to\bullet}.\, (t\, p) \land (t\, \overline{p}) \land (H\, \lambda x^\bullet.\, t(\langle a \rangle x)) \land (H\, \lambda x^\bullet.\, t([a]x)) \\ \land \big(H\, \lambda x_1^\bullet.\, (H\, \lambda x_2^\bullet.\, t\,(x_1 \lor x_2))\big) \land \big(H\, \lambda x_1^\bullet.\, (H\, \lambda x_2^\bullet.\, t\,(x_1 \land x_2))\big),$$

where w.l.o.g. we restrict attention to a single atomic proposition p and action a. For any given $\chi^{\bullet\to\bullet}$, one can show that $H\,\chi \equiv \bigcap_{\psi\in\text{ML}}(\chi\,\psi)$. For example, consider $\psi = \langle a\rangle p$ and any $\chi^{\bullet\to\bullet}$. Using fixpoint unfolding and β-reduction, we obtain that $H\,\chi \equiv \chi\, p \land \cdots \land (H\, \lambda x^\bullet.\, \chi(\langle a\rangle x)) \land \cdots$, which further expands to $\chi\, p \land \cdots \land \big(\chi\, \langle a\rangle p \land \cdots\big) \land \cdots$.

Coming back to GFP^1_φ, the macro H is used in the clause $H\,(\lambda z^\bullet.F\, z \to (\lambda x^\bullet.\varphi(F))\, z)$. By the above characterisation, this is equivalent to $\bigcap_{\psi\in\text{ML}}(\lambda z^\bullet.F\, z \to (\lambda x^\bullet.\varphi(F))\, z)\,\psi$, which tests for every ML-definable set, represented by the defining formula, whether the fixpoint has stabilised at this set. With Lem. 2 we then obtain the following.

Theorem 3. *The* HFL^1*-alternation hierarchy collapses to the alternation-free fragment in* HFL^2 *over* $\mathbb{T}^1_{\mathsf{fin}}$.

Note that the collapse is not achieved by replacing HFL^1 formulas of the form $\nu F^{\bullet\to\bullet}.\varphi$ by *semantically equivalent* formulas in HFL^2 using at most least fixpoints. The HFL^2 formula GFP^1_φ defines a function of type $\bullet \to \bullet$ and uses only least fixpoints, but it agrees with $\nu F^{\bullet\to\bullet}.\varphi$ on all HFL^1-definable arguments, not necessarily on arbitrary ones. This, however, is sufficient to obtain Thm. 3 above.

Finally, the pattern of H can be extended to

$$H_2 = \nu H_2^{((\bullet\to\bullet)\to\bullet)\to\bullet}.\, \lambda t^{(\bullet\to\bullet)\to\bullet}.\, (t\,(\lambda x^\bullet.\, p)) \land (t\,(\lambda x^\bullet.\, \overline{p})) \land (t\,(\lambda x^\bullet.\, x)) \\ \land \big(H_2\, \lambda f^{\bullet\to\bullet}.\, t(\lambda x^\bullet.\, \langle a\rangle (f\, x))\big) \land \big(H_2\, \lambda f^{\bullet\to\bullet}.\, t(\lambda x^\bullet.\, [a](f\, x))\big) \\ \land \big(H_2\, \lambda f_1^{\bullet\to\bullet}.\, (H_2\, \lambda x_2^{\bullet\to\bullet}.\, t\,(\lambda x^\bullet.\, (f_1\, x) \lor f_2\, x))\big) \land \big(H_2\, \lambda f_1^{\bullet\to\bullet}.\, (H_2\, \lambda x_2^{\bullet\to\bullet}.\, t\,(\lambda x^\bullet.\, (f_1\, x) \land f_2\, x))\big),$$

which enumerates all HFL^1 formulas of the form $\lambda x^\bullet.\,\varphi$ where φ is an ML formula that may also contain x. Similar reasoning as above yields that for a collapse result, it is enough to test for fixpoint stabilisation on these formulas. Using

$$\mathsf{GFP}^2_\varphi := \Big(\mu g^\tau.\lambda F^{\tau'}\lambda f^{(\bullet\to\bullet}.\big(F\, f \land H_2\, (\lambda f'^{\bullet\to\bullet}.F\, f' \to (\lambda f^{\bullet\to\bullet}.\varphi(F))\, f')\big) \lor (g\,\varphi(F))\, f\Big)\; \lambda f^{(\bullet\to\bullet)}.\top$$

with $\tau' = ((\bullet \to \bullet) \to \bullet)$ and $\tau = \tau' \to (\bullet \to \bullet) \to \bullet$, we obtain the following result.

Theorem 4. *The* HFL^2*-alternation hierarchy collapses to the alternation-free fragment in* HFL^3 *over* $\mathbb{T}^2_{\mathsf{fin}}$.

3 Discussion

Thms. 1, 3 and 4 extend results for $\mathcal{L}_\mu$ that yield a collapse of the alternation hierarchy over special classes of structures. The picture for the entire HFL is much farther from completion: we have strictness in HFL^k for $k \leq 2$ over the class of all structures but a collapse for HFL^k inside HFL^{k+1} over the classes $\mathbb{T}^k_{\mathsf{fin}}$, for $k \in \{0, 1, 2\}$. For $k \geq 3$, the set of formulas to test on for stabilisation is not necessarily free of λ abstractions, and established results on the number of variables used in λ expressions of type order 3 prohibit such an extension without further inspection of the set of formulas to test on.

It should be noted that a hierarchy collapse over the classes $\mathbb{T}^k_{\mathsf{fin}}$ is not the same as the (trivial) collapse over a single structure with finite closure ordinals (where every formula is equivalent to a fixpoint free formula that defines the same set). The translations presented here yield alternation-free formulas that are uniformly equivalent over the full classes $\mathbb{T}^0_{\mathsf{fin}}$,$\mathbb{T}^1_{\mathsf{fin}}$ and $\mathbb{T}^2_{\mathsf{fin}}$.

One can construct non-trivial infinite structures in $\mathbb{T}^0_{\mathsf{fin}} \setminus \mathbb{T}^{\sim}_{\mathsf{fin}}$; it appears that $\mathbb{T}^k_{\mathsf{fin}}$ is a rather rich class of structures beyond finite structures. We suspect in fact that all the inclusions stated in Eq. 1 are strict. A formal proof would advance the current understanding of the expressive power of type-order fragments of HFL significantly and therefore add to the understanding of the interplay between program logics and higher-order features.

At last, it is also not clear how fixpoint alternation behaves over the class of all structures above type level 2. Moreover, optimality of the constructions here is also open; it is still possible that the alternation hierarchy for HFL^k already collapses inside HFL^k for certain k and $\mathbb{T}^k_{\mathsf{fin}}$ or a subclass thereof.

References

[1] 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.

[2] J. C. Bradfield. The modal μ-calculus alternation hierarchy is strict. In *CONCUR'96*, volume 1119 of *LNCS*, pages 233–246. Springer, 1996.

[3] F. Bruse. Alternation is strict for higher-order modal fixpoint logic. In *GandALF'16*, volume 226 of *EPTCS*, pages 105–119, 2016.

[4] J. Gutierrez, F. Klaedtke, and M. Lange. The μ-calculus alternation hierarchy collapses over structures with restricted connectivity. *Theoretical Computer Science*, 560(3):292–306, 2014.

[5] R. Kaivola. Axiomatising linear time mu-calculus. In *CONCUR'95*, volume 962 of *LNCS*, pages 423–437. Springer, 1995.

[6] N. Kobayashi, É. Lozes, and F. Bruse. On the relationship between higher-order recursion schemes and higher-order fixpoint logic. In *POPL'17*, pages 246–259. ACM, 2017.

[7] D. Kozen. Results on the propositional μ-calculus. *TCS*, 27:333–354, December 1983.

[8] E. Lozes. A type-directed negation elimination. In *FICS'15*, volume 191 of *EPTCS*, pages 132–142, 2015.

[9] M. Müller-Olm. A modal fixpoint logic with chop. In *STACS'99*, volume 1563 of *LNCS*, pages 510–520. Springer, 1999.

[10] M. Viswanathan and R. Viswanathan. A higher order modal fixed point logic. In *CONCUR'04*, volume 3170 of *LNCS*, pages 512–528. Springer, 2004.