Global Flattening of Nested Inductive Definitions

Paul Blain Levy
School of Computer Science, University of Birmingham, Edgbaston, Birmingham, UK
p.b.levy@bham.ac.uk

For a calculus of nested inductive definitions, interpreted over posets using least prefixpoints, we see how to regard the semantics of closed terms as a single mutually inductive definition.

By duality, the same result applies to a calculus of nested coinductive definitions.

1 Introduction

The purpose of this work is to investigate "global flattening": a rather simple property of nested fixpoint calculations that appears not to have been previously noted in the literature. The property says that, for a calculus of inductive definitions, the semantics of closed terms can be given by a single mutually inductive definition.

In this short abstract, we shall treat only posets, omitting the categorical version, in which least prefixpoints are replaced by initial algebras.

We first recall some standard notions of fixpoint theory in Section 2, and then explain the idea via an example calculus in Section 3. Our general result is stated and proved in Section 4, and Section 6 considers coinductive definitions.

2 Order-theoretic preliminaries

For any set A and endofunction $f : A \to A$, an element $x \in A$ is *f-fixed* or an *f-fixpoint* when $f(x) = x$. We may omit f when the context makes it clear.

For any poset A and endofunction f that is monotone (i.e. order-preserving), an element $x \in A$ is *f-prefixed* or an *f-prefixpoint* when $f(x) \leqslant x$, and dually *f-postfixed* or an *f-postfixpoint* when $x \leqslant f(x)$. [1] Thus being fixed is equivalent to being both prefixed and postfixed. Any supremum of prefixpoints is prefixed, and dually any infimum of postfixpoints is postfixed.

The least f-prefixpoint, if it exists, is written μf and said to be "inductively defined". It is necessarily a fixpoint—this is called the "inductive inversion principle". Dually the greatest postfixpoint, if it exists, is necessarily a fixpoint; it is written νf and said to be "coinductively defined".

An *ω-cpo* is a poset where every ω-chain $x_0 \leqslant x_1 \leqslant \cdots$ has a supremum. A map between ω-cpos is *ω-continuous* when it is monotone and preserves all these suprema.

An *ω-cppo* is an ω-cpo with a least element, written $\bot$. On an ω-cppo A, every ω-continuous endofunction f has a least prefixpoint, given by $\mu f = \bigvee_{n \in \mathbb{N}} f^n(\bot)$.

There are other well-known ways to ensure that μf exists. For example, given a monotone endofunction f on a complete lattice, we obtain μf as the infimum of all prefixpoints, and dually νf as the supremum of all postfixpoints. (This is a version of Tarski's fixpoint theorem.)

[1] This terminology follows [6]. Some authors use the opposite terminology, following [5].

Submitted to:
FICS 2026

3 Example calculus

We illustrate global flattening by taking a (somewhat arbitrary) example of an ω-cppo equipped with two elements and an ω-continuous binary operation. The ω-cppo is actually a complete lattice: it is $\mathscr{P}\mathbb{N}$, ordered by inclusion. The two elements are the set of all even natural numbers, written Even, and the set of all primes, written Prime. The binary operation $*$ is given by

$$A * B \stackrel{\text{def}}{=} \{x^2 + 3xyz + 51 \mid x,y \in A, z \in B\}$$

Now let us formulate a calculus. The syntax is given by

$$M \stackrel{\text{def}}{=} \mathtt{Even} \mid \mathtt{Prime} \mid M * N \mid x \mid \mu x.M$$

where x ranges over variables. As usual, we write $\Gamma \vdash M$, where Γ is a context (repetition-free list of variables), to say that M is a term whose free variables all appear in Γ. This judgement is defined inductively in the usual way. We identify α-equivalent terms, i.e. terms that have the same operation symbols, variables and binding structure but may differ in the choice of bound variables. We write Closed for the set of all closed terms.

The semantics of the calculus is organized as follows. Firstly, for a context Γ, a Γ-*environment* ρ provides for each variable $x \in \Gamma$ an element $\rho_x \in \mathscr{P}\mathbb{N}$, and we write $[\![\Gamma]\!]$ for the set of all such. Each term $\Gamma \vdash M$ will denote an ω-continuous function from $[\![\Gamma]\!]$ to $\mathscr{P}\mathbb{N}$, written $[\![\Gamma \vdash M]\!]$, or $[\![M]\!]$ for short. These denotations are defined by recursion on the judgement $\Gamma \vdash M$ in the usual way, and we then formulate and prove the substitution lemma in the usual way. Details are in Section 4.

Writing ε for the empty environment, we now see that each closed term $M \in \mathsf{Closed}$ has a denotation $[\![M]\!]\varepsilon \in \mathscr{P}\mathbb{N}$. So we have an element $([\![M]\!]\varepsilon)_{M \in \mathsf{Closed}}$ of the complete lattice $(\mathscr{P}\mathbb{N})^{\mathsf{Closed}}$, and our task is to describe it as a least prefixpoint. Accordingly, we define the ω-continuous endofunction H on $(\mathscr{P}\mathbb{N})^{\mathsf{Closed}}$, sending d to the family

$$\begin{array}{rcl} \mathtt{Even} & \mapsto & \mathsf{Even} \\ \mathtt{Prime} & \mapsto & \mathsf{Prime} \\ M * N & \mapsto & d(M) * d(N) \\ \mu x.M & \mapsto & d(M[\mu x.M/x]) \end{array}$$

While $([\![M]\!]\varepsilon)_{M \in \mathsf{Closed}}$ is easily seen to be H-fixed (via the substitution lemma), we shall show that it is actually the least H-prefixpoint. We call such a result "global flattening" because all the nested inductive definitions appearing in terms are replaced by a single mutually inductive definition, taken across all closed terms.

4 The result for a many-sorted signature

In Section 3, our calculus had two constant symbols and one binary operation symbol. To generalize this, recall that a *single-sorted signature* $\mathscr{S} = (\mathsf{Ar}(f))_{f \in \mathsf{Op}}$ is a set Op of *operation symbols*, where each $f \in \mathsf{Op}$ is equipped with a set called its *arity*, written $\mathsf{Ar}(f)$.

More generally still, let T be a set of *sorts*. Then a T-*sorted signature* $\mathscr{S} = (f : (\mathsf{In}^i_f)_{i \in \mathsf{Ar}(f)} \to \mathsf{Out}_f)_{f \in \mathsf{Op}}$ is a set Op of *operation symbols*, where each $f \in \mathsf{Op}$ is equipped with an indexed family of sorts $(\mathsf{In}^i_f)_{i \in \mathsf{Ar}(f)}$ and a sort Out_f. The pair $(T, \mathscr{S})$ is called a *many-sorted signature*, and will be fixed for the rest of the paper.

Now we define the calculus. A context Γ is a list of declarations $x : A$, where A is a sort, with no variable appearing more than once. Our typing judgement is $\Gamma \vdash M : B$, where Γ is a context and B a sort. The typing rules are as follows. For each operation symbol f of arity I, we still have the rule

$$\frac{(\Gamma \vdash M_i : A_i)_{i\in I}}{\Gamma \vdash f(M_i)_{i\in I} : B}\ f : (A_i)_{i\in I} \to B$$

For each sort A we have the following rules:

$$\frac{}{\Gamma \vdash x : A}\ (x : A) \in \Gamma \qquad \frac{\Gamma, x : A \vdash M : A}{\Gamma \vdash \mu x.M : A}$$

For posets A and B, a *partial monotone function* $f : A \rightharpoonup B$ is a partial function that is monotone on its domain. A *partial element* of a poset A is either an element or undefined. Here are some constructions:

- Given a partial function $f : \prod_{i\in I} A_i \rightharpoonup B$ and partial element x of A_i for all $i \in I$, we obtain a partial element $f(x_i)_{i\in I}$ of B. It is undefined when either the family $(x_i)_{i\in I}$ is not total, or it is total but not in $\mathsf{dom}(f)$.
- Given partial monotone endofunction f on a poset A, we obtain a partial element μf. It is undefined when either f is not total or it is total but has no least prefixpoint.

An *interpretation* for $(T, \mathscr{S})$ associates to each sort A a poset $[\![A]\!]$, and to each operation symbol f a (total) monotone function $\overline{f} : \prod_{i\in \mathsf{Ar}(f)} [\![\mathsf{In}^i_f]\!] \to [\![\mathsf{Out}_f]\!]$. Such an interpretation yields a semantics of the calculus in the usual way. Firstly, for a context Γ, we write $[\![\Gamma]\!] \stackrel{\text{def}}{=} \prod_{(x:A)\in\Gamma} [\![A]\!]$ for the set of all Γ-environments. Each term $\Gamma \vdash M : A$ will denote a partial monotone function $[\![M]\!] : [\![\Gamma]\!] \to [\![A]\!]$, and these denotations are defined recursively as follows:

$$\begin{array}{lcccl} [\![f(M_i)_{i\in I}]\!] & : & \rho & \mapsto & \overline{f}([\![M_i]\!]\rho)_{i\in I} \\ [\![x]\!] & : & \rho & \mapsto & \rho_x \\ [\![\mu x.M]\!] & : & \rho & \mapsto & \mu a.\,[\![M]\!](\rho, x \mapsto a) \end{array}$$

A closed term M is *semantically defined* when $[\![M]\!]\varepsilon$ is defined.

For contexts Γ and Δ, a *substitution* $k : \Gamma \to \Delta$ sends each variable $(x : A) \in \Gamma$ to a term $\Delta \vdash k_x : A$. Given this and a term $\Gamma \vdash M : B$, we obtain the term $\Delta \vdash M[k] : A$ from M by replacing each variable x by k_x. The *substitution lemma* says that the partial function $[\![M[k]]\!]$ extends the partial function $\rho \mapsto [\![M]\!](x \mapsto [\![k_x]\!]\rho)$. This is proved using a renaming lemma.

A set $\mathscr{L}$ of closed terms is *closed*[2] when

- $f(M_i)_{i\in I} \in \mathscr{L}$ and $i \in I$ implies $M_i \in \mathscr{L}$
- $\mu x.M \in \mathscr{L}$ implies $M[\mu x.M/x] \in \mathscr{L}$.

For example, the set of all semantically defined closed terms. Now let $\mathscr{L}$ be any closed set of semantically defined closed terms. Thus we have an element $([\![M]\!]\varepsilon)_{M\in\mathscr{L}}$ of the poset $\prod_{(M:A)\in\mathscr{L}} [\![A]\!]$, and our task is to describe it as a least prefixpoint. By analogy with Section 3, define the monotone endofunction $H_{\mathscr{L}}$ on $\prod_{(M:A)\in\mathscr{L}} [\![A]\!]$, sending d to the family

$$\begin{array}{rcl} f(M_i)_{i\in I} & \mapsto & \overline{f}(d(M_i))_{i\in I} \\ \mu x.M & \mapsto & d(M[\mu x.M/x]) \end{array}$$

Now we give our key result:

[2] Kozen [4] noted that this coincides with the Fischer-Ladner notion of closure [2] via the translation of PDL into μ-calculus.

Theorem 1 (Global flattening). *The family* $([\![M]\!]\varepsilon)_{M\in\mathscr{L}}$ *is the least prefixpoint of* $H_{\mathscr{L}}$.

Proof. The substitution lemma shows that the family $([\![M]\!]\varepsilon)_{M\in\mathscr{L}}$ is $H_{\mathscr{L}}$-fixed, so it only remains to prove leastness.

First we introduce some terminology. For any context Γ, say that a *semantic* Γ*-environment* is an element of $[\![\Gamma]\!]$, and an $\mathscr{L}$*-syntactic* Γ*-environment* sends each $(x : A) \in \Gamma$ to a closed term $\sigma_x \in \mathscr{L}$ of type A. Each element $d \in \prod_{(M:A)\in\mathscr{L}} [\![A]\!]$ induces a map d^Γ from $\mathscr{L}$-syntactic to semantic Γ-environments, sending σ to $(d(\sigma_x))_{(x:A)\in\Gamma}$.

Let d be an $H_{\mathscr{L}}$-prefixpoint, which means that

$$\begin{array}{rcl} \overline{f}(d(M_i))_{i\in I} & \leqslant & d(f(M_i)_{i\in I}) \\ d(M[\mu x.M/x]) & \leqslant & d(\mu x.M) \end{array}$$

We show inductively on the judgement $\Gamma \vdash M : B$ that, for every $\mathscr{L}$-syntactic Γ-environment σ such that $M[\sigma] \in \mathscr{L}$ and $[\![M]\!]d^\Gamma(\sigma)$ exists, we have

$$[\![M]\!]d^\Gamma(\sigma) \leqslant d(M[\sigma])$$

□

5 Nested mutually inductive definitions

So far, our calculus provides single inductive definitions but not inductively defined pairs or ω-sequences. To incorporate these, suppose we have not only a many-sorted signature $(T, \mathscr{S})$ but also a set Cmpd of *compounds*, where each compound r represents a family of sorts written $Decode(r)$.

A typing context Γ now consists of declarations of the form $x : r$, where r is a compound. As before, the typing judgement takes the form $\Gamma \vdash t : A$, where A is a sort. The typing rules are as follows. For each operation symbol f of arity I, we still have the rule

$$\frac{(\Gamma \vdash M_i : A_i)_{i\in I}}{\Gamma \vdash f(M_i)_{i\in I} : B}\; f : (A_i)_{i\in I} \to B$$

For each compound r with $Decode(r) = (A_i)_{i\in I}$ and each $\hat{\imath} \in I$, we have the rules

$$\frac{}{\Gamma \vdash x\,\hat{\imath} : A_{\hat{\imath}}}\; (x : r) \in \Gamma \qquad \frac{(\Gamma, x : r \vdash M_i : A_i)_{i\in I}}{\Gamma \vdash \mu x.\,(M_i)_{i\in I}\,\hat{\imath} : A}$$

The semantics goes through as before, deeming $w\,\hat{\imath}$ to be undefined whenever w is undefined.

Here are two cases of special interest.

- Let Cmpd be the set of sorts, each representing the singleton of itself. Then we simply obtain the calculus of Section 4.
- Let Cmpd be the set of lists of sort, with each list $A_0, \ldots, A_{n-1}$ representing the family $(A_i)_{i<n}$. Then we obtain the so-called "vectorial" calculus, allowing inductively defined n-tuples for $n \in \mathbb{N}$ but not inductively defined ω-sequences.

To adapt the story from Section 4, we have to distinguish between two notions. A set of closed terms $\mathscr{L}$ is *closed* when it has the following properties:

- $f(M_i)_{i\in I} \in \mathscr{L}$ and $i \in I$ implies $M_i \in \mathscr{L}$

- $\mu x.\,(M_i)_{i \in I}\,\hat{\imath} \in \mathscr{L}$ implies that $M_{\hat{\imath}}[\mu x.M/x] \in \mathscr{L}$

It is *laterally closed* when it additionally satisfies the following:

- $\mu x.\,(M_i)_{i \in I}\,\hat{\imath} \in \mathscr{L}$ implies that, for all $j \in I$, we have $\mu x.\,(M_i)_{i \in I}\,j \in \mathscr{L}$

It turns out that Theorem 1 holds for any laterally closed set $\mathscr{L}$ of semantically defined closed terms. The proof is essentially unchanged.

6 Coinductive definitions

By duality, global flattening holds also for a calculus of nested *coinductive* definitions. However, it remains to be seen whether it can be adapted to a mixed calculus of inductive and coinductive definitions. For this purpose, the fixpoint construction given in [3] may be helpful.

7 Related work

Let us now look at some related results.

First is Bekić's theorem, which expresses an inductively defined pair $\mu(x,y).\,(f(x,y),g(x,y))$ in terms of single inductive definitions. This generalizes to an inductive defined n-tuple, for any natural number n. (However, it does not work for infinite collections.) So any term in the vectorial calculus can be represented in the non-vectorial calculus. This is the opposite of flattening.

Next are results such as

$$\mu x.\,\mu y.\,f(x,y) = \mu z.,f(z,z) \tag{1}$$

These can be called "local flattening" results, as they apply to a single point in a term.

Lastly, we consider [1, Theorem 2.7.19], which gives (for any finitary signature) a normal form for terms with k levels of alternation. The $k = 1$ case, in conjunction with the equation (1), says that any term M with only μ's can be reduced to a vectorial term of the form $(\mu x.\,N)0$. This is also a consequence of Theorem 1, by taking $\mathscr{L}$ to be the least closed set containing M, which is finite.

Despite this point of similarity between the two lines of work, there are significant differences:

- Most importantly, [1, Theorem 2.7.19] treats terms that mix μ and ν, whereas our work does not.
- Our work allows infinitary operation symbols and infinite compounds, whereas [1] works in the finitary setting.
- Our result gives the semantics of all closed terms via a single inductive definition, whereas [1, Theorem 2.7.19] treats each closed term separately.
- Our Theorem 1 gives an explicit and simple description of the endofunction used, whereas in [1, Theorem 2.7.19] the endofunction's description is more complicated and relegated to the proof.

References

[1] André Arnold & Damian Niwiński (2001): *Rudiments of μ-calculus*. Studies in Logic and the Foundations of Mathematics 146, Elsevier.

[2] Michael J. Fischer & Richard E. Ladner (1979): *Propositional dynamic logic of regular programs*. Journal of Computer and System Sciences 18(2), pp. 194–211, doi:https://doi.org/10.1016/0022-0000(79)90046-1. Available at `https://www.sciencedirect.com/science/article/pii/0022000079900461`.

[3] Charles Grellois & Paul-André Melliès (2015): *An Infinitary Model of Linear Logic*. In Andrew Pitts, editor: *Foundations of Software Science and Computation Structures*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 41–55, doi:10.1007/978-3-662-46678-0_3.

[4] Dexter Kozen (1983): *Results on the propositional μ-calculus*. *Theoretical Computer Science* 27(3), pp. 333–354, doi:https://doi.org/10.1016/0304-3975(82)90125-6. Available at `https://www.sciencedirect.com/science/article/pii/0304397582901256`. Special Issue Ninth International Colloquium on Automata, Languages and Programming (ICALP) Aarhus, Summer 1982.

[5] Zohar Manna & Adi Shamir (1978): *The convergence of functions to fixedpoints of recursive definitions*. *Theoretical Computer Science* 6(2), pp. 109–141, doi:10.1016/0304-3975(78)90033-6.

[6] Michael Smyth & Gordon D. Plotkin (1982): *The category-theoretic solution of recursive domain equations*. *SIAM Journal on Computing* 11, pp. 761–783, doi:10.1137/0211062.