

Saturated Semantics for Coalgebraic Logic Programming

Filippo Bonchi and Fabio Zanasi

ENS Lyon, U. de Lyon, CNRS, INRIA, UCBL, France

Abstract. A series of recent papers introduces a coalgebraic semantics for logic programming, where the behavior of a goal is represented by a *parallel* model of computation called coinductive tree. This semantics fails to be compositional, in the sense that the coalgebra formalizing such behavior does not commute with the substitutions that may apply to a goal. We suggest that this is an instance of a more general phenomenon, occurring in the setting of interactive systems (in particular, nominal process calculi), when one tries to model their semantics with coalgebras on presheaves. In those cases, compositionality can be obtained through *saturation*. We apply the same approach to logic programming: the resulting semantics is compositional and enjoys an elegant formulation in terms of coalgebras on presheaves and their right Kan extensions.

1 Introduction

Coalgebras on presheaves have been successfully employed to provide semantics to *nominal* calculi: sophisticated process calculi with complex mechanisms for variable binding, like the π -calculus [?,?]. The idea is to have an index category \mathbf{C} of interfaces (or names), and encode as a presheaf $\mathcal{F}: \mathbf{C} \rightarrow \mathbf{Set}$ the mapping of any object i of \mathbf{C} to the set of states having i as interface, and any arrow $f: i \rightarrow j$ to a function switching the interface of states from i to j . The operational semantics of the calculus will arise as a notion of transition between states, that is, as a coalgebra $\alpha: \mathcal{F} \rightarrow \mathcal{B}(\mathcal{F})$, where $\mathcal{B}: \mathbf{Set}^{\mathbf{C}} \rightarrow \mathbf{Set}^{\mathbf{C}}$ is a functor on presheaves encoding the kind of behavior that we want to express.

As an arrow in a presheaf category, α has to be a natural transformation, i.e. it should commute with arrows $f: i \rightarrow j$ in the index category \mathbf{C} . Unfortunately, this naturality requirement may fail when the structure of \mathbf{C} is rich enough, as for instance when non-injective substitutions [?,?] or name fusions [?,?] occur. As a concrete example, consider the π -calculus term $t = \bar{a}\langle x \rangle | b(y)$ consisting of a process $\bar{a}\langle x \rangle$ sending a message x on a channel named a , in parallel with $b(y)$ receiving a message on a channel named b . Since the names a and b are different, the two processes cannot synchronize. Conversely the term $t\theta = \bar{a}\langle x \rangle | a(y)$, that is obtained by applying the substitution θ mapping b to a , can synchronize. If θ is an arrow of the index category \mathbf{C} , then the operational semantics α is not natural since $\alpha(t\theta) \neq \alpha(t)\bar{\theta}$, where $\bar{\theta}$ denotes the application of θ to the transitions of t . As a direct consequence, also the unique morphism to the

terminal coalgebra is not natural: this means that the abstract semantics of π -calculus is not *compositional* - in other words, bisimilarity is not a congruence w.r.t. name substitutions. In order to make bisimilarity a congruence, Sangiorgi introduced in [?] *open bisimilarity*, that is defined by considering the transitions of processes under *all* possible name substitutions θ .

The approach of *saturated semantics* [?] can be seen as a generalization of open bisimilarity, relying on analogous principles: the operational semantics α is “saturated” w.r.t. the arrows of the index category \mathbf{C} , resulting in a natural transformation α^\sharp in $\mathbf{Set}^{\mathbf{C}}$. In [?,?], this is achieved by first shifting the definition of α to the category $\mathbf{Set}^{|\mathbf{C}|}$ of presheaves indexed by the discretization $|\mathbf{C}|$ of \mathbf{C} . Since $|\mathbf{C}|$ does not have other arrow than the identities, α is trivially a natural transformation in this setting. The source of α is $\mathcal{U}(\mathcal{F}) \in \mathbf{Set}^{|\mathbf{C}|}$, where $\mathcal{U}: \mathbf{Set}^{\mathbf{C}} \rightarrow \mathbf{Set}^{|\mathbf{C}|}$ is a forgetful functor defined by composition with the inclusion $\iota: |\mathbf{C}| \rightarrow \mathbf{C}$. The functor \mathcal{U} has a right adjoint $\mathcal{K}: \mathbf{Set}^{|\mathbf{C}|} \rightarrow \mathbf{Set}^{\mathbf{C}}$ sending a presheaf to its *right Kan extension* along ι . The adjoint pair $\mathcal{U} \dashv \mathcal{K}$ induces an isomorphism $(-)_X^\sharp: \mathbf{Set}^{|\mathbf{C}|}[\mathcal{U}(X), Y] \rightarrow \mathbf{Set}^{\mathbf{C}}[X, \mathcal{K}(Y)]$ mapping α to α^\sharp . The latter is a natural transformation in $\mathbf{Set}^{\mathbf{C}}$ and, consequently, the abstract semantics results to be compositional.

In this paper, we show that the saturated approach can be fruitfully instantiated to *coalgebraic logic programming* [?,?,?], which consists of a novel semantics for logic programming and a parallel resolution algorithm based on *coinductive trees*. These are a variant of and-or trees [?] modeling *parallel* implementations of logic programming, where the soundness of the derivations represented by a tree is guaranteed by the restriction to *term-matching* (whose algorithm, differently from unification, is parallelizable [?]).

There are two analogies with the π -calculus: (a) the state space is modeled by a presheaf on the index category \mathbf{L}_Σ^{op} , that is the (opposite) *Lawvere Theory* associated with some signature Σ ; (b) the operational semantics given in [?] fails to be a natural transformation in $\mathbf{Set}^{\mathbf{L}_\Sigma^{op}}$: Example ?? provides a counterexample which is similar to the π -calculus term t discussed above.

The authors of [?] obviate to (b) by relaxing naturality to *lax naturality*: the operational semantics p of a logic program is given as an arrow in the category $\mathbf{Lax}(\mathbf{L}_\Sigma^{op}, \mathbf{Poset})$ of locally ordered functors $\mathcal{F}: \mathbf{L}_\Sigma^{op} \rightarrow \mathbf{Poset}$ and lax natural transformations between them. They show the existence of a cofree comonad that induces a morphism $\llbracket - \rrbracket_p$ mapping atoms (i.e., atomic formulae) to coinductive trees. Since $\llbracket - \rrbracket_p$ is not natural but lax natural, the semantics provided by coinductive trees is not compositional, in the sense that, for some atoms A and substitution θ ,

$$\llbracket A\theta \rrbracket_p \neq \llbracket A \rrbracket_p \bar{\theta}$$

where $\llbracket A\theta \rrbracket_p$ is the coinductive tree associated with $A\theta$ and $\llbracket A \rrbracket_p \bar{\theta}$ denotes the result of applying θ to each atom occurring in the tree $\llbracket A \rrbracket_p$.

Instead of introducing laxness, we propose to tackle the non-naturality of p with a saturated approach. It turns out that, in the context of logic programming, the saturation map $(-)_\sharp$ has a neat description in terms of substitution

mechanisms: while p performs *term-matching* between the atoms and the heads of clauses of a given logic program, its saturation p^\sharp (given as a morphism in $\mathbf{Set}^{\mathbf{L}_\Sigma^{op}}$) performs *unification*. It is worth to remark here that not only most general unifiers are considered but *all* possible unifiers.

A cofree construction leading to a map $\llbracket - \rrbracket_{p^\sharp}$ can be obtained by very standard categorical tools, such as terminal sequences [?]. This is possible because, as \mathbf{Set} , both $\mathbf{Set}^{\mathbf{L}_\Sigma^{op}}$ and $\mathbf{Set}^{|\mathbf{L}_\Sigma^{op}|}$ are (co)complete categories, whereas in the lax approach, $\mathit{Lax}(\mathbf{L}_\Sigma^{op}, \mathbf{Poset})$ not being (co)complete, more indirect and more sophisticated categorical constructions are needed [?, Sec. 4]. By naturality of p^\sharp , the semantics given by $\llbracket - \rrbracket_{p^\sharp}$ turns out to be compositional, as in the desiderata. Analogously to $\llbracket - \rrbracket_p$, also $\llbracket - \rrbracket_{p^\sharp}$ maps atoms to tree structures, which we call *saturated trees*. They generalize coinductive trees, in the sense that the latter can be seen as a “desaturation” of saturated trees, where all unifiers that are not term-matchers have been discarded. This observation leads to a *translation* from saturated to coinductive trees, based on the counit ϵ of the adjunction $\mathcal{U} \dashv \mathcal{K}$. It follows that our framework encompasses the semantics in [?,?].

Analogously to what is done in [?], we propose a notion of *refutation subtree* of a given saturated tree, intuitively corresponding to an SLD-refutation of an atomic goal in a program. This leads to a result of soundness and completeness of our semantics with respect to SLD-resolution, crucially using both compositionality and the translation into coinductive trees.

Related works. Apart from [?,?,?], there exist other categorical perspectives on (extensions of) logic programming, such as [?,?,?]. Amongst these, the most relevant for us is [?] since it exploits a form of saturation: states representing formulae are both instantiated by substitution and contextualized by other formulae in “and”. Beyond logic programming, the idea of exploiting saturation to achieve compositionality is even older than [?]. As far as we know, [?] is the first work where saturation is explored in terms of coalgebras. It is interesting to note that, in [?], a subset of the same authors also proposed laxness as a solution for the lack of compositionality of Petri nets.

A third approach, alternative to laxness and saturation, may be possible by taking a special kind of “powerobject” functor as done in [?,?] for giving a coalgebraic semantics to fusion and open π -calculus. We have chosen saturated semantics for its generality: it works for any behavioral functor \mathcal{B} and it models a phenomenon that occurs in many different computational models (see e.g. [?]).

Acknowledgements. We thank E. Komendantskaya, T. Hirschowitz, D. Petrisan, J. Power, M. Sammartino, the *Plume* team and the anonymous referees for the helpful comments. Our work is supported by project ANR 12IS02001 PACE.

2 Coalgebraic Logic Programming

In this section we recall the framework of coalgebraic logic programming, as introduced in [?,?,?]. For this purpose, we first fix some terminology and notation, mainly concerning category theory and logic programming.

Given a (small) category \mathbf{C} , $|\mathbf{C}|$ denotes the category with the same objects as \mathbf{C} but no other arrow than the identities. With a little abuse of notation, $o \in |\mathbf{C}|$ indicates that o is an object of \mathbf{C} and $\mathbf{C}[o_1, o_2]$ the set of arrows from o_1 to o_2 . A \mathbf{C} -indexed *presheaf* is any functor $\mathcal{G}: \mathbf{C} \rightarrow \mathbf{Set}$. We write $\mathbf{Set}^{\mathbf{C}}$ for the category of \mathbf{C} -indexed presheaves and natural transformations between them. Given a functor $\mathcal{B}: \mathbf{C} \rightarrow \mathbf{C}$, a \mathcal{B} -*coalgebra* on $o \in |\mathbf{C}|$ is an arrow $p: o \rightarrow \mathcal{B}(o)$.

We fix a *signature* Σ of function symbols, each equipped with a fixed arity, and a countably infinite set $Var = \{x_1, x_2, x_3, \dots\}$ of variables. We model substitutions and unification of terms over Σ and Var according to the categorical perspective of [?,?]. To this aim, let the (opposite) *Lawvere Theory* of Σ be a category \mathbf{L}_{Σ}^{op} where objects are natural numbers, with $n \in |\mathbf{L}_{\Sigma}^{op}|$ intuitively representing variables x_1, x_2, \dots, x_n from Var . For any two $n, m \in |\mathbf{L}_{\Sigma}^{op}|$, the set $\mathbf{L}_{\Sigma}^{op}[n, m]$ consists of all n -tuples $\langle t_1, \dots, t_n \rangle$ of terms where only variables among x_1, \dots, x_m occur. The identity on $n \in |\mathbf{L}_{\Sigma}^{op}|$, denoted by id_n , is given by the tuple $\langle x_1, \dots, x_n \rangle$. The composition of $\langle t_1^1, \dots, t_n^1 \rangle: n \rightarrow m$ and $\langle t_1^2, \dots, t_m^2 \rangle: m \rightarrow m'$ is the tuple $\langle t_1, \dots, t_n \rangle: n \rightarrow m'$, where t_i is the term t_i^1 in which every variable x_j has been replaced with t_j^2 , for $1 \leq j \leq m$ and $1 \leq i \leq n$.

We call *substitutions* the arrows of \mathbf{L}_{Σ}^{op} and use Greek letters θ, σ and τ to denote them. Given $\theta_1: n \rightarrow m_1$ and $\theta_2: n \rightarrow m_2$, a *unifier* of θ_1 and θ_2 is a pair of substitutions $\sigma: m_1 \rightarrow m$ and $\tau: m_2 \rightarrow m$, where m is some object of \mathbf{L}_{Σ}^{op} , such that $\sigma \circ \theta_1 = \tau \circ \theta_2$. The *most general unifier* of θ_1 and θ_2 is a unifier with a universal property, i.e. a pushout of the diagram $m_1 \xleftarrow{\theta_1} n \xrightarrow{\theta_2} m_2$.

An *alphabet* \mathcal{A} consists of a signature Σ , a set of variables Var and a set of predicate symbols P, P_1, P_2, \dots each assigned an arity. Given P of arity n and Σ -terms t_1, \dots, t_n , $P(t_1, \dots, t_n)$ is called an *atom*. We use Latin capital letters A, B, \dots for atoms. Given a substitution $\theta = \langle t_1, \dots, t_n \rangle: n \rightarrow m$ and an atom A with variables among x_1, \dots, x_n , we adopt the standard notation of logic programming in denoting with $A\theta$ the atom obtained by replacing x_i with t_i in A , for $1 \leq i \leq n$. The atom $A\theta$ is called a *substitution instance* of A . The notation $\{A_1, \dots, A_m\}\theta$ is a shorthand for $\{A_1\theta, \dots, A_m\theta\}$. Given atoms A_1 and A_2 , we say that A_1 *unifies* with A_2 (equivalently, they are *unifiable*) if they are of the form $A_1 = P(t_1, \dots, t_n)$, $A_2 = P(t'_1, \dots, t'_n)$ and a unifier $\langle \sigma, \tau \rangle$ of $\langle t_1, \dots, t_n \rangle$ and $\langle t'_1, \dots, t'_n \rangle$ exists. Observe that, by definition of unifier, this amounts to saying that $A_1\sigma = A_2\tau$. *Term matching* is a particular case of unification, where σ is the identity substitution. In this case we say that $\langle \sigma, \tau \rangle$ is a *term-matcher* of A_1 and A_2 , meaning that $A_1 = A_2\tau$.

A *logic program* \mathbb{P} consists of a finite set of *clauses* C written as $H \leftarrow B_1, \dots, B_k$. The components H and B_1, \dots, B_k are atoms, where H is called the *head* of C and B_1, \dots, B_k form the *body* of C . One can think of $H \leftarrow B_1, \dots, B_k$ as representing the first-order formula $(B_1 \wedge \dots \wedge B_k) \rightarrow H$. We say that \mathbb{P} is *ground* if only ground atoms (i.e. without variables) occur in its clauses. The central algorithm of logic programming is SLD-resolution, checking whether a finite set of atoms (called a *goal*) is *refutable* in \mathbb{P} and giving a substitution called *computed answer* as output. Relevant for our exposition are *and-or trees* [?], which represent executions of SLD-resolution exploiting two forms of parallelism: *and-*

parallelism, corresponding to simultaneous refutation-search of multiple atoms in a goal, and *or-parallelism*, exploring multiple attempts to refute the same goal.

Definition 1. *Given a logic program \mathbb{P} and an atom A , the (parallel) and-or tree for A in \mathbb{P} is the possibly infinite tree T satisfying the following properties:*

1. *Each node in T is either an and-node or an or-node.*
2. *Each and-node is labeled with one atom and its children are or-nodes.*
3. *The root of T is an and-node labeled with A .*
4. *Each or-node is labeled with \bullet and its children are and-nodes.*
5. *For every and-node s in T , let A' be its label. For every clause $H \leftarrow B_1, \dots, B_k$ of \mathbb{P} and most general unifier $\langle \sigma, \tau \rangle$ of A' and H , s has exactly one child t , and viceversa. For each atom B in $\{B_1, \dots, B_k\}\tau$, t has exactly one child labeled with B , and viceversa.*

As standard for any tree, we have a notion of *depth*: the root is at depth 0 and depth $i + 1$ is given by the children of nodes at depth i .

2.1 The Ground Case

We recall the coalgebraic semantics of ground logic programs introduced in [?]. For the sequel we fix an alphabet \mathcal{A} , a set At of ground atoms and a ground logic program \mathbb{P} . The behavior of \mathbb{P} is represented by a coalgebra $p: At \rightarrow \mathcal{P}_f \mathcal{P}_f(At)$ on **Set**, where \mathcal{P}_f is the finite powerset functor and p is defined as follows:

$$p: A \mapsto \{\{B_1, \dots, B_k\} \mid H \leftarrow B_1, \dots, B_k \text{ is a clause of } \mathbb{P} \text{ and } A = H\}.$$

The idea is that p maps an atom $A \in At$ to the set of bodies of clauses of \mathbb{P} whose head H unifies with A , i.e. (in the ground case) $A = H$. Therefore $p(A) \in \mathcal{P}_f \mathcal{P}_f(At)$ can be seen as representing the and-or tree of A in \mathbb{P} up to depth 2, according to Definition ???: each element $\{B_1, \dots, B_k\}$ of $p(A)$ corresponds to a child of the root, whose children are labeled with B_1, \dots, B_k . The full tree is recovered as an element of $\mathcal{C}(\mathcal{P}_f \mathcal{P}_f)(At)$, where $\mathcal{C}(\mathcal{P}_f \mathcal{P}_f)$ is the *cofree comonad* on $\mathcal{P}_f \mathcal{P}_f$, standardly provided by the following construction [?,?].

Construction 1 *The terminal sequence for the functor $At \times \mathcal{P}_f \mathcal{P}_f(-): \mathbf{Set} \rightarrow \mathbf{Set}$ consists of sequences of objects X_α and arrows $\delta_\alpha: X_{\alpha+1} \rightarrow X_\alpha$, defined by induction on α as follows.*

$$X_\alpha := \begin{cases} At & \alpha = 0 \\ At \times \mathcal{P}_f \mathcal{P}_f(X_\beta) & \alpha = \beta + 1 \end{cases} \quad \delta_\alpha := \begin{cases} \pi_1 & \alpha = 0 \\ id_{At} \times \mathcal{P}_f \mathcal{P}_f(\delta_\beta) & \alpha = \beta + 1 \end{cases}$$

For α a limit ordinal, X_α is given as a limit of the sequence and a function $\delta_\alpha: X_\alpha \rightarrow X_\beta$ is given for each $\beta < \alpha$ by the limiting property of X_α .

By [?] it follows that the sequence given above converges to a limit X_γ such that $X_\gamma \cong X_{\gamma+1}$. Since $X_{\gamma+1}$ is defined as $At \times \mathcal{P}_f \mathcal{P}_f(X_\gamma)$, there is a projection function $\pi_2: X_{\gamma+1} \rightarrow \mathcal{P}_f \mathcal{P}_f(X_\gamma)$ which makes $\pi_2 \circ \delta_\gamma^{-1}: X_\gamma \rightarrow \mathcal{P}_f \mathcal{P}_f(X_\gamma)$ the cofree $\mathcal{P}_f \mathcal{P}_f$ -coalgebra on At . This induces the cofree comonad $\mathcal{C}(\mathcal{P}_f \mathcal{P}_f): \mathbf{Set} \rightarrow \mathbf{Set}$ on $\mathcal{P}_f \mathcal{P}_f$ as a functor mapping At to X_γ .

As the elements of the cofree comonad on \mathcal{P}_f are standardly presented as finitely branching trees [?], those for $\mathcal{P}_f\mathcal{P}_f$ can be seen as finitely branching trees with two sorts of nodes occurring at alternating depth. We now define a $\mathcal{C}(\mathcal{P}_f\mathcal{P}_f)$ -coalgebra $\llbracket - \rrbracket_p: At \rightarrow \mathcal{C}(\mathcal{P}_f\mathcal{P}_f)(At)$.

Construction 2 Given a ground program \mathbb{P} , let $p: At \rightarrow \mathcal{P}_f\mathcal{P}_f(At)$ be the coalgebra associated with \mathbb{P} . We define a cone $\{p_\alpha: At \rightarrow X_\alpha\}_{\alpha < \gamma}$ on the terminal sequence of Construction ?? as follows:

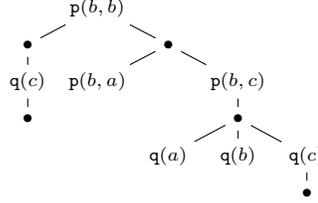
$$p_\alpha := \begin{cases} id_{At} & \alpha = 0 \\ \langle id_{At}, (\mathcal{P}_f\mathcal{P}_f(p_\beta) \circ p) \rangle & \alpha = \beta + 1. \end{cases}$$

For α a limit ordinal, $p_\alpha: At \rightarrow X_\alpha$ is provided by the limiting property of X_α . Then in particular $X_\gamma = \mathcal{C}(\mathcal{P}_f\mathcal{P}_f)(At)$ yields a function $\llbracket - \rrbracket_p: At \rightarrow \mathcal{C}(\mathcal{P}_f\mathcal{P}_f)(At)$.

Given an atom $A \in At$, the tree $\llbracket A \rrbracket_p \in \mathcal{C}(\mathcal{P}_f\mathcal{P}_f)(At)$ is built by iteratively applying the map p , first to A , then to each atom in $p(A)$, and so on. For each natural number m , p_m maps A to its and-or tree up to depth m . As shown in [?], the limit $\llbracket - \rrbracket_p$ of all such approximations provides the full and-or tree of A .

Example 1. Consider the ground logic program on the left-hand side, based on an alphabet consisting of a signature $\{a^0, b^0, c^0\}$ and predicates $\mathbf{p}(-, -)$, $\mathbf{q}(-)$. The and-or tree $\llbracket \mathbf{p}(b, b) \rrbracket_p \in \mathcal{C}(\mathcal{P}_f\mathcal{P}_f)(At)$ is depicted on the right-hand side.

$$\begin{aligned} \mathbf{p}(b, c) &\leftarrow \mathbf{q}(a), \mathbf{q}(b), \mathbf{q}(c) \\ \mathbf{p}(b, b) &\leftarrow \mathbf{p}(b, a), \mathbf{p}(b, c) \\ \mathbf{p}(b, b) &\leftarrow \mathbf{q}(c) \\ \mathbf{q}(c) &\leftarrow \end{aligned}$$



2.2 The General Case

We recall the extension of the coalgebraic semantics to arbitrary (i.e. possibly non-ground) logic programs presented in [?,?]. In presence of variables, and-or trees are not guaranteed to represent sound derivations, whence *coinductive trees* are introduced as a sound variant of and-or trees, where unification is restricted to term-matching. We refer to [?,?] for more details.

Before formally defining coinductive trees, it is worth recalling that, in [?], the collection of atoms (based on an alphabet \mathcal{A}) is modeled as a presheaf $At: \mathbf{L}_\Sigma^{op} \rightarrow \mathbf{Set}$. The index category is the (opposite) *Lawvere Theory* \mathbf{L}_Σ^{op} of Σ , as defined above. For each natural number $n \in |\mathbf{L}_\Sigma^{op}|$, $At(n)$ is defined as the set of atoms with variables among x_1, \dots, x_n . Given an arrow $\theta \in \mathbf{L}_\Sigma^{op}[n, m]$, the function $At(\theta): At(n) \rightarrow At(m)$ is defined by substitution, i.e. $At(\theta)(A) := A\theta$. By definition, whenever an atom A belongs to $At(n)$, then it also belongs to $At(n')$, for all $n' \geq n$. However, the occurrences of the same atom in $At(n)$ and $At(n')$ (for $n \neq n'$) are considered distinct: the atoms $A \in At(n)$ and $A \in At(n')$ can be thought of as two states $x_1, \dots, x_n \vdash A$ and $x_1, \dots, x_{n'} \vdash A$ with two different interfaces x_1, \dots, x_n and $x_1, \dots, x_{n'}$. For this reason, when referring to an atom A , it is important to always specify the set $At(n)$ to which it belongs.

Definition 2. Given a logic program \mathbb{P} , a natural number n and an atom $A \in \text{At}(n)$, the n -coinductive tree for A in \mathbb{P} is the possibly infinite tree T satisfying properties 1-4 of Definition ?? and property 5 replaced by the following¹:

5. For every and-node s in T , let $A' \in \text{At}(n)$ be its label. For every clause $H \leftarrow B_1, \dots, B_k$ of \mathbb{P} and term-matcher $\langle id_n, \tau \rangle$ of A' and H , with $B_1\tau, \dots, B_k\tau \in \text{At}(n)$, s has exactly one child t , and viceversa. For each atom B in $\{B_1, \dots, B_k\}\tau$, t has exactly one child labeled with B , and viceversa.

We recall from [?] the categorical formalization of this class of trees. The first step is to generalize the definition of the coalgebra p associated with a program \mathbb{P} . Definition ?? suggests how p should act on an atom $A \in \text{At}(n)$, for a fixed n :

$$A \mapsto \{ \{B_1, \dots, B_k\}\tau \mid H \leftarrow B_1, \dots, B_k \text{ is a clause of } \mathbb{P}, \\ A = H\tau \text{ and } B_1\tau, \dots, B_k\tau \in \text{At}(n) \}. \quad (1)$$

For each clause $H \leftarrow B_1, \dots, B_k$, there might be infinitely (but countably) many substitutions τ such that $A = H\tau$ (see e.g. [?]). Thus the object on the right-hand side of (??) will be associated with the functor $\mathcal{P}_c\mathcal{P}_f: \mathbf{Set} \rightarrow \mathbf{Set}$, where \mathcal{P}_c and \mathcal{P}_f are respectively the countable powerset functor and the finite powerset functor. In order to formalize this as a coalgebra on $\text{At}: \mathbf{L}_\Sigma^{op} \rightarrow \mathbf{Set}$, consider liftings $\widetilde{\mathcal{P}}_c: \mathbf{Set}^{\mathbf{L}_\Sigma^{op}} \rightarrow \mathbf{Set}^{\mathbf{L}_\Sigma^{op}}$ and $\widetilde{\mathcal{P}}_f: \mathbf{Set}^{\mathbf{L}_\Sigma^{op}} \rightarrow \mathbf{Set}^{\mathbf{L}_\Sigma^{op}}$, standardly defined on presheaves $\mathcal{F}: \mathbf{L}_\Sigma^{op} \rightarrow \mathbf{Set}$ by postcomposition respectively with \mathcal{P}_c and \mathcal{P}_f . Then one would like to fix (??) as the definition of the n -component of a natural transformation $p: \text{At} \rightarrow \widetilde{\mathcal{P}}_c\widetilde{\mathcal{P}}_f(\text{At})$. The key problem with this formulation is that p would *not* be a natural transformation, as shown by the following example.

Example 2. Consider the signature $\Sigma = \{cons^2, succ^1, zero^0, nil^0\}$ and the predicates $\text{List}(-)$, $\text{Nat}(-)$. The program NatList , encoding the definition of lists of natural numbers, will be our running example of a non-ground logic program.

$$\begin{array}{ll} \text{List}(cons(x_1, x_2)) \leftarrow \text{Nat}(x_1), \text{List}(x_2) & \text{List}(nil) \leftarrow \\ \text{Nat}(succ(x_1)) \leftarrow \text{Nat}(x_1) & \text{Nat}(zero) \leftarrow \end{array}$$

Fix a substitution $\theta = \langle nil \rangle: 1 \rightarrow 0$ and, for each $n \in |\mathbf{L}_\Sigma^{op}|$, suppose that $p(n): \text{At}(n) \rightarrow \widetilde{\mathcal{P}}_c\widetilde{\mathcal{P}}_f(\text{At})(n)$ is defined according to (??). Then the square

$$\begin{array}{ccc} \text{At}(1) & \xrightarrow{p(1)} & \widetilde{\mathcal{P}}_c\widetilde{\mathcal{P}}_f(\text{At})(1) \\ \text{At}(\theta) \downarrow & & \downarrow \widetilde{\mathcal{P}}_c\widetilde{\mathcal{P}}_f(\text{At})(\theta) \\ \text{At}(0) & \xrightarrow{p(0)} & \widetilde{\mathcal{P}}_c\widetilde{\mathcal{P}}_f(\text{At})(0) \end{array}$$

¹ Our notion of coinductive tree corresponds to the notion of coinductive forest of breadth n as in [?, Def.4.4], the only difference being that we “glue” together all trees of the forest into a single tree.

does not commute. A counterexample is provided by the atom $\mathbf{List}(x_1) \in At(1)$. Passing through the bottom-left corner of the square, $\mathbf{List}(x_1)$ is mapped first to $\mathbf{List}(nil) \in At(0)$ and then to $\{\emptyset\} \in \widetilde{\mathcal{P}}_c \widetilde{\mathcal{P}}_f(At)(0)$ - intuitively, this yields a refutation of the goal $\{\mathbf{List}(x_1)\}$ with substitution of x_1 with nil . Passing through the top-right corner, $\mathbf{List}(x_1)$ is mapped first to $\emptyset \in \widetilde{\mathcal{P}}_c \widetilde{\mathcal{P}}_f(At)(1)$ and then to $\emptyset \in \widetilde{\mathcal{P}}_c \widetilde{\mathcal{P}}_f(At)(0)$, i.e. the computation ends up in a failure.

In [?, Sec.4] the authors overcome this difficulty by relaxing the naturality requirement. The morphism p is defined as a $\check{\mathcal{P}}_c \check{\mathcal{P}}_f$ -coalgebra in the category $Lax(\mathbf{L}_\Sigma^{op}, \mathbf{Poset})$ of locally ordered functors $\mathcal{F}: \mathbf{L}_\Sigma^{op} \rightarrow \mathbf{Poset}$ and *lax* natural transformations, with each component $p(n)$ given according to (??) and $\check{\mathcal{P}}_c \check{\mathcal{P}}_f$ the extension of $\widetilde{\mathcal{P}}_c \widetilde{\mathcal{P}}_f$ to an endofunctor on $Lax(\mathbf{L}_\Sigma^{op}, \mathbf{Poset})$.

The lax approach fixes the problem, but presents also some drawbacks. Unlike the categories \mathbf{Set} and $\mathbf{Set}^{\mathbf{L}_\Sigma^{op}}$, $Lax(\mathbf{L}_\Sigma^{op}, \mathbf{Poset})$ is neither complete nor cocomplete, meaning that a cofree comonad on $\check{\mathcal{P}}_c \check{\mathcal{P}}_f$ cannot be retrieved through the standard Constructions ?? and ?? that were used in the ground case. Moreover, the category of $\check{\mathcal{P}}_c \check{\mathcal{P}}_f$ -coalgebrae becomes problematic, because coalgebra maps are subject to a commutativity property stricter than the one of lax natural transformations. These two issues force the formalization of non-ground logic program to use quite different (and more sophisticated) categorical tools than the ones employed for the ground case. Finally, as stressed in the Introduction, the laxness of p makes the resulting semantics not compositional.

3 Saturated Semantics

Motivated by the observations of the previous section, we propose a *saturated approach* to the semantics of logic programs. For this purpose, we consider an adjunction between presheaf categories as depicted on the left.

$$\begin{array}{ccc}
 \mathbf{Set}^{\mathbf{L}_\Sigma^{op}} & \begin{array}{c} \xrightarrow{\mathcal{U}} \\ \perp \\ \xleftarrow{\mathcal{K}} \end{array} & \mathbf{Set}^{|\mathbf{L}_\Sigma^{op}|} \\
 & & \\
 & & \begin{array}{ccc} |\mathbf{L}_\Sigma^{op}| & \xhookrightarrow{\iota} & \mathbf{L}_\Sigma^{op} \\ \mathcal{F} \downarrow & \swarrow \mathcal{K}(\mathcal{F}) & \\ \mathbf{Set} & & \end{array}
 \end{array}$$

The left adjoint \mathcal{U} is the forgetful functor, given by precomposition with the inclusion functor $\iota: |\mathbf{L}_\Sigma^{op}| \hookrightarrow \mathbf{L}_\Sigma^{op}$. As shown in [?, Th.X.1], \mathcal{U} has a right adjoint $\mathcal{K}: \mathbf{Set}^{|\mathbf{L}_\Sigma^{op}|} \rightarrow \mathbf{Set}^{\mathbf{L}_\Sigma^{op}}$ sending $\mathcal{F}: |\mathbf{L}_\Sigma^{op}| \rightarrow \mathbf{Set}$ to its *right Kan extension* along ι . This is a presheaf $\mathcal{K}(\mathcal{F}): \mathbf{L}_\Sigma^{op} \rightarrow \mathbf{Set}$ mapping an object n of \mathbf{L}_Σ^{op} to

$$\mathcal{K}(\mathcal{F})(n) := \prod_{\theta \in \mathbf{L}_\Sigma^{op}[n, m]} \mathcal{F}(m)$$

where m is any object of \mathbf{L}_Σ^{op} . Intuitively, $\mathcal{K}(\mathcal{F})(n)$ is a set of tuples indexed by arrows with source n and such that, at index $\theta: n \rightarrow m$, there are elements of

$\mathcal{F}(m)$. We use $\dot{x} \dot{y}, \dots$ to denote such tuples and we write $\dot{x}(\theta)$ to denote the element at index θ of the tuple \dot{x} . Alternatively, when it is important to show how the elements depend from the indexes, we use $\langle x \rangle_{\theta:n \rightarrow m}$ (or simply $\langle x \rangle_\theta$) to denote the tuple having at index θ the element x . With this notation, we can express the behavior of $\mathcal{K}(\mathcal{F}): \mathbf{L}_\Sigma^{op} \rightarrow \mathbf{Set}$ on an arrow $\theta: n \rightarrow m$ as

$$\mathcal{K}(\mathcal{F})(\theta): \dot{x} \mapsto \langle \dot{x}(\sigma \circ \theta) \rangle_{\sigma:m \rightarrow m'}. \quad (2)$$

The tuple $\langle \dot{x}(\sigma \circ \theta) \rangle_\sigma$ in $\mathcal{K}(\mathcal{F})(m)$ can be intuitively read as follows: for each $\sigma \in \mathbf{L}_\Sigma^{op}[m, m']$, we let the element indexed by σ be the one which was indexed by $\sigma \circ \theta \in \mathbf{L}_\Sigma^{op}[n, m']$ in the input tuple \dot{x} .

All this concerns the behavior of \mathcal{K} on the objects of $\mathbf{Set}^{|\mathbf{L}_\Sigma^{op}|}$. For an arrow $f: \mathcal{F} \rightarrow \mathcal{G}$ in $\mathbf{Set}^{|\mathbf{L}_\Sigma^{op}|}$, the natural transformation $\mathcal{K}(f)$ is defined as an indexwise application of f on tuples from $\mathcal{K}(\mathcal{F})$. For all $n \in |\mathbf{L}_\Sigma^{op}|$, $\dot{x} \in \mathcal{K}(\mathcal{F})(n)$,

$$\mathcal{K}(f)(n): \dot{x} \mapsto \langle f(m)(\dot{x}(\theta)) \rangle_{\theta:n \rightarrow m}.$$

For any presheaf $\mathcal{F}: \mathbf{L}_\Sigma^{op} \rightarrow \mathbf{Set}$, the unit η of the adjunction is instantiated to a morphism $\eta_{\mathcal{F}}: \mathcal{F} \rightarrow \mathcal{KU}(\mathcal{F})$ given as follows: for all $n \in |\mathbf{L}_\Sigma^{op}|$, $X \in \mathcal{F}(n)$,

$$\eta_{\mathcal{F}}(n): X \mapsto \langle \mathcal{F}(\theta)(X) \rangle_{\theta:n \rightarrow m}.$$

When taking \mathcal{F} to be At , $\eta_{At}: At \rightarrow \mathcal{KU}(At)$ maps an atom to its *saturation*: for each $A \in At(n)$, the tuple $\eta_{At}(n)(A)$ consists of all substitution instances $At(\theta)(A) = A\theta$ of A , each indexed by the corresponding $\theta \in \mathbf{L}_\Sigma^{op}[n, m]$.

As shown in Example ??, given a program \mathbb{P} , the family of functions p defined by (??) fails to be a morphism in $\mathbf{Set}^{|\mathbf{L}_\Sigma^{op}|}$. However, it forms a morphism in $\mathbf{Set}^{|\mathbf{L}_\Sigma^{op}|}$

$$p: \mathcal{U}At \rightarrow \widehat{\mathcal{P}}_c \widehat{\mathcal{P}}_f(\mathcal{U}At)$$

where $\widehat{\mathcal{P}}_c$ and $\widehat{\mathcal{P}}_f$ denote the liftings of \mathcal{P}_c and \mathcal{P}_f to $\mathbf{Set}^{|\mathbf{L}_\Sigma^{op}|}$. The naturality requirement is trivially satisfied in $\mathbf{Set}^{|\mathbf{L}_\Sigma^{op}|}$, since $|\mathbf{L}_\Sigma^{op}|$ is discrete. The adjunction induces a morphism $p^\sharp: At \rightarrow \mathcal{K}\widehat{\mathcal{P}}_c \widehat{\mathcal{P}}_f \mathcal{U}(At)$ in $\mathbf{Set}^{|\mathbf{L}_\Sigma^{op}|}$, defined as

$$At \xrightarrow{\eta_{At}} \mathcal{KU}(At) \xrightarrow{\mathcal{K}(p)} \mathcal{K}\widehat{\mathcal{P}}_c \widehat{\mathcal{P}}_f \mathcal{U}(At). \quad (3)$$

In the sequel, we write \mathcal{S} for $\mathcal{K}\widehat{\mathcal{P}}_c \widehat{\mathcal{P}}_f \mathcal{U}$. The idea is to let \mathcal{S} play the same role as $\mathcal{P}_f \mathcal{P}_f$ in the ground case, with the coalgebra $p^\sharp: At \rightarrow \mathcal{S}(At)$ encoding the program \mathbb{P} . An atom $A \in At(n)$ is mapped to $\langle p(m)(A\sigma) \rangle_{\sigma:n \rightarrow m}$, that is:

$$p^\sharp(n): A \mapsto \langle \{ \{ B_1, \dots, B_k \} \tau \mid H \leftarrow B_1, \dots, B_k \text{ is a clause of } \mathbb{P}, \\ A\sigma = H\tau \text{ and } B_1\tau, \dots, B_k\tau \in At(m) \} \rangle_{\sigma:n \rightarrow m}. \quad (4)$$

Intuitively, $p^\sharp(n)$ retrieves all unifiers $\langle \sigma, \tau \rangle$ of A and heads of \mathbb{P} : first, $A\sigma \in At(m)$ arises as a component of the saturation of A , according to $\eta_{At}(n)$; then, the substitution τ is given by term-matching on $A\sigma$, according to $\mathcal{K}(p)(m)$.

By naturality of p^\sharp , we achieve the property of ‘‘commuting with substitutions’’ that was precluded by the term-matching approach, as shown by the following rephrasing of Example ??.

Example 3. Consider the same square of Example ??, with p^\sharp in place of p and \mathcal{S} in place of $\widetilde{\mathcal{P}}_c\widetilde{\mathcal{P}}_f$. The atom $\mathbf{List}(x_1) \in At(1)$ together with the substitution $\theta = \langle nil \rangle : 1 \rightarrow 0$ does not constitute a counterexample to commutativity anymore. Indeed $p^\sharp(1)$ maps $\mathbf{List}(x_1)$ to the tuple $\langle p(n)(\mathbf{List}(x_1)\sigma) \rangle_{\sigma: 1 \rightarrow n}$, which is then mapped by $\mathcal{S}(At)(\theta)$ to $\langle p(n)(\mathbf{List}(x_1)\sigma' \circ \theta) \rangle_{\sigma': 0 \rightarrow n}$ according to (?). Observe that the latter is just the tuple $\langle p(n)(\mathbf{List}(nil)\sigma') \rangle_{\sigma': 0 \rightarrow n}$ obtained by applying first $At(\theta)$ and then $p^\sharp(0)$ to the atom $\mathbf{List}(x_1)$.

Another benefit of saturated semantics is that $p^\sharp: At \rightarrow \mathcal{S}(At)$ lives in a (co)complete category which behaves (pointwise) as **Set**. This allows us to follow the same steps as in the ground case, constructing a coalgebra for the cofree comonad $\mathcal{C}(\mathcal{S})$ as a straightforward generalization of Constructions ?? and ??.

Construction 3 *The terminal sequence for the functor $At \times \mathcal{S}(-)$: $\mathbf{Set}^{\mathbf{L}^{\mathcal{S}}} \rightarrow \mathbf{Set}^{\mathbf{L}^{\mathcal{S}}}$ consists of a sequence of objects X_α and arrows $\delta_\alpha: X_{\alpha+1} \rightarrow X_\alpha$, which are defined just as in Construction ??, with \mathcal{S} replacing $\mathcal{P}_f\mathcal{P}_f$. By using [?, Th.7], it can be checked that this sequence converges to a limit X_γ such that $X_\gamma \cong X_{\gamma+1}$ and X_γ is the carrier of the cofree \mathcal{S} -coalgebra on At .*

Since \mathcal{S} is accessible, the cofree comonad $\mathcal{C}(\mathcal{S})$ exists and maps At to X_γ given as in Construction ?. A $\mathcal{C}(\mathcal{S})$ -coalgebra $\llbracket - \rrbracket_{p^\sharp}: At \rightarrow \mathcal{C}(\mathcal{S})(At)$ is given below.

Construction 4 *The terminal sequence for $At \times \mathcal{S}(-)$ induces a cone $\{p_\alpha^\sharp: At \rightarrow X_\alpha\}_{\alpha < \gamma}$ as in Construction ?? with p^\sharp and \mathcal{S} replacing p and $\mathcal{P}_f\mathcal{P}_f$. This yields a natural transformation $\llbracket - \rrbracket_{p^\sharp}: At \rightarrow X_\gamma$, where $X_\gamma = \mathcal{C}(\mathcal{S})(At)$.*

As in the ground case, the coalgebra $\llbracket - \rrbracket_{p^\sharp}$ is constructed as an iterative application of p^\sharp : we call *saturated tree* the associated tree structure.

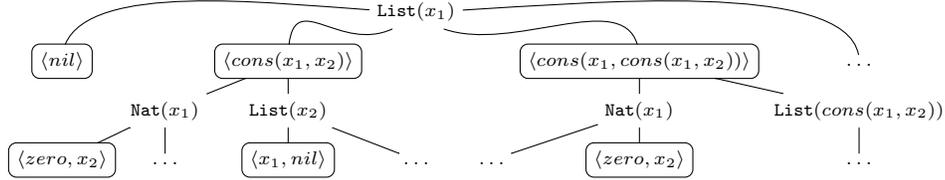
Definition 3. *Given a logic program \mathbb{P} , a natural number n and an atom $A \in At(n)$, the saturated tree for A in \mathbb{P} is the possibly infinite tree T satisfying properties 1-3 of Definition ?? and properties 4 and 5 replaced by the following:*

4. *Each or-node is labeled with a substitution σ and its children are and-nodes.*
5. *For every and-node s in T , let $A' \in At(n')$ be its label. For every clause $H \leftarrow B_1, \dots, B_k$ of \mathbb{P} and unifier $\langle \sigma, \tau \rangle$ of A' and H , with $\sigma: n' \rightarrow m'$ and $B_1\tau, \dots, B_k\tau \in At(m')$, s has exactly one child t labeled with σ , and viceversa. For each atom B in $\{B_1, \dots, B_k\}\tau$, t has exactly one child labeled with B , and viceversa.*

We have now seen three kinds of tree, exhibiting different substitution mechanisms. In saturated trees one considers all the unifiers, whereas in and-or trees and coinductive trees one restricts to most general unifiers and term-matchers respectively. Moreover, in a coinductive tree each and-node is labeled with an atom in $At(n)$ for a fixed n , while in a saturated tree n can dynamically change.

Example 4. Part of the infinite saturated tree of $\mathbf{List}(x_1) \in At(1)$ in $\mathbf{NatList}$ is depicted below. Note that not all labels of and-nodes belong to $At(1)$, as it would

be the case for a coinductive tree: such information is inherited from the label of the parent or-node, which is now a substitution. For instance, both $\mathbf{Nat}(x_1)$ and $\mathbf{List}(x_2)$ belong to $At(2)$, since their parent is labeled with $\langle cons(x_1, x_2) \rangle: 1 \rightarrow 2$ (using the convention that the target of a substitution is the largest index appearing among its variables).



We can generalize these observations to the following adequacy theorem.

Theorem 1. *Let $\llbracket - \rrbracket_{p^\#}$ be defined from a program \mathbb{P} according to Construction ???. Then, for all n and $A \in At(n)$, the saturated tree of A in \mathbb{P} is $\llbracket A \rrbracket_{p^\#}$.*

In the above theorem and in the rest of the paper, with an abuse of notation we use $\llbracket A \rrbracket_{p^\#}$ to denote the application of $\llbracket - \rrbracket_{p^\#}(n)$ to $A \in At(n)$ without mentioning the object $n \in |\mathbf{L}_\Sigma^{op}|$. For an arrow $\theta \in \mathbf{L}_\Sigma^{op}[n, m]$, we write $\bar{\theta}$ for $\mathcal{C}(\mathcal{S})(At)(\theta): \mathcal{C}(\mathcal{S})(At)(n) \rightarrow \mathcal{C}(\mathcal{S})(At)(m)$. With this notation, we can state the following theorem that is an immediate consequence of the naturality of $\llbracket - \rrbracket_{p^\#}$.

Theorem 2 (Compositionality). *For all atoms $A \in At(n)$ and substitutions $\theta \in \mathbf{L}_\Sigma^{op}[n, m]$,*

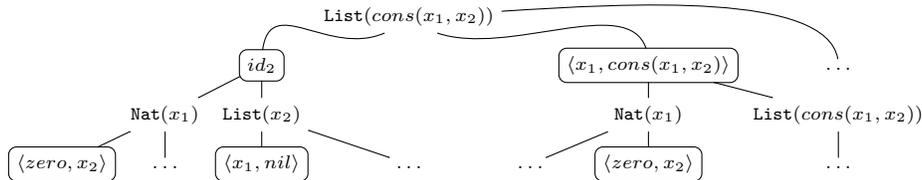
$$\llbracket A\theta \rrbracket_{p^\#} = \llbracket A \rrbracket_{p^\#} \bar{\theta}.$$

We conclude this section with a concrete description of the behavior of the operator $\bar{\theta}$, for a given substitution $\theta \in \mathbf{L}_\Sigma^{op}[n, m]$. Let r be the root of a tree $T \in \mathcal{C}(\mathcal{S})(At)(n)$ and r' the root of $T\bar{\theta}$. Then

1. the node r has label A iff r' has label $A\theta$;
2. the node r has a child t with label $\sigma \circ \theta$ and children t_1, \dots, t_n iff r' has a child t' with label σ and children $t_1 \dots t_n$.

Note that the children t_1, \dots, t_n are exactly the same in both trees: $\bar{\theta}$ only modifies the root and the or-nodes at depth 1 of T , while it leaves untouched all the others. This peculiar behavior can be better understood by observing that the definition of $\mathcal{K}(\mathcal{F})(\theta)$, as in (??), is independent of the presheaf \mathcal{F} . As a result, $\bar{\theta} = X_\gamma(\theta)$ is independent of all the X_α s built in Construction ??.

Example 5. Recall from Example ?? the saturated tree $\llbracket \mathbf{List}(x_1) \rrbracket_{p^\#}$. For $\theta = \langle cons(x_1, x_2) \rangle$, the tree $\llbracket \mathbf{List}(x_1) \rrbracket_{p^\#} \bar{\theta}$ is depicted below.



4 Desaturation

One of the main features of coinductive trees is to represent (sound) and-or parallel derivations of goals. This leads the authors of [?] to a resolution algorithm exploiting the two forms of parallelism. Motivated by these developments, we include coinductive trees in our framework, showing how they can be obtained as a “desaturation” of saturated trees.

For this purpose, the key ingredient is given by the counit ϵ of the adjunction $\mathcal{U} \dashv \mathcal{K}$. Given a presheaf $\mathcal{F}: |\mathbf{L}_\Sigma^{op}| \rightarrow \mathbf{Set}$, the morphism $\epsilon_{\mathcal{F}}: \mathcal{U}\mathcal{K}(\mathcal{F}) \rightarrow \mathcal{F}$ is defined as follows: for all $n \in |\mathbf{L}_\Sigma^{op}|$ and $\dot{x} \in \mathcal{U}\mathcal{K}(\mathcal{F})(n)$,

$$\epsilon_{\mathcal{F}}(n): \dot{x} \mapsto \dot{x}(id_n) \quad (5)$$

where $\dot{x}(id_n)$ is the element of the input tuple \dot{x} which is indexed by the identity substitution $id_n \in \mathbf{L}_\Sigma^{op}[n, n]$. In the logic programming perspective, the intuition is that, while the unit of the adjunction provides the saturation of an atom, the counit reverses the process. It takes the saturation of an atom and gives back the substitution instance given by the identity, that is, the atom itself.

The next construction defines a morphism $\bar{d}: \mathcal{U}(\mathcal{C}(\mathcal{S})(At)) \rightarrow \mathcal{C}(\widehat{\mathcal{P}}_c \widehat{\mathcal{P}}_f)(\mathcal{U}At)$ where $\mathcal{C}(\widehat{\mathcal{P}}_c \widehat{\mathcal{P}}_f): \mathbf{Set}^{|\mathbf{L}_\Sigma^{op}|} \rightarrow \mathbf{Set}^{|\mathbf{L}_\Sigma^{op}|}$ is the cofree comonad on $\widehat{\mathcal{P}}_c \widehat{\mathcal{P}}_f$, obtained through a terminal sequence analogously to Construction ???. The idea is that \bar{d} acts on saturated trees as the depthwise application of $\epsilon_{\mathcal{U}At}$.

Construction 5 For α an ordinal, let us note by Y_α the objects occurring in the construction of $\mathcal{C}(\widehat{\mathcal{P}}_c \widehat{\mathcal{P}}_f)(\mathcal{U}At)$ and with X_α the ones in the construction of $\mathcal{C}(\mathcal{S})(At)$, converging to $X_\gamma = \mathcal{C}(\mathcal{S})(At)$. We define a sequence $\{d_\alpha: \mathcal{U}(X_\alpha) \rightarrow Y_\alpha\}_{\alpha < \gamma}$ in $\mathbf{Set}^{|\mathbf{L}_\Sigma^{op}|}$ as follows:

$$d_\alpha := \begin{cases} id_{\mathcal{U}At} & \alpha = 0 \\ id_{\mathcal{U}At} \times (\widehat{\mathcal{P}}_c \widehat{\mathcal{P}}_f(d_\beta) \circ \epsilon_{\widehat{\mathcal{P}}_c \widehat{\mathcal{P}}_f \mathcal{U}(X_\beta)}) & \alpha = \beta + 1. \end{cases}$$

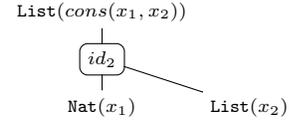
For $\alpha < \gamma$ a limit ordinal, $d_\alpha: \mathcal{U}(X_\alpha) \rightarrow Y_\alpha$ is provided by the limiting property of Y_α . This sequence induces a morphism $\bar{d}: \mathcal{U}(\mathcal{C}(\mathcal{S})(At)) \rightarrow \mathcal{C}(\widehat{\mathcal{P}}_c \widehat{\mathcal{P}}_f)(\mathcal{U}At)$.

The next theorem shows that \bar{d} is a translation from saturated to coinductive trees: given an atom $A \in At(n)$, it maps $\llbracket A \rrbracket_{p^\#}$ to the n -coinductive tree of A . The key intuition is that n -coinductive trees can be seen as saturated trees where the labeling of or-nodes has been restricted to the identity substitution id_n , represented as \bullet (see Definition ??). The operation of pruning all or-nodes (and their descendants) in $\llbracket A \rrbracket_{p^\#}$ which are not labeled with id_n is precisely what is provided by Construction ??, in virtue of the definition of the counit ϵ given in (??).

Theorem 3 (Desaturation). Let $\llbracket - \rrbracket_{p^\#}: At \rightarrow \mathcal{C}(\mathcal{S})(At)$ be defined for a logic program \mathbb{P} according to Construction ?? and $\bar{d}: \mathcal{U}(\mathcal{C}(\mathcal{S})(At)) \rightarrow \mathcal{C}(\widehat{\mathcal{P}}_c \widehat{\mathcal{P}}_f)(\mathcal{U}At)$ be defined according to Construction ??. Then for all $n \in |\mathbf{L}_\Sigma^{op}|$ and $A \in \mathcal{U}At(n)$, the n -coinductive tree of A in \mathbb{P} is $(\bar{d} \circ \mathcal{U}(\llbracket - \rrbracket_{p^\#})) (n)(A)$.

Theorem ?? also provides an alternative formalization for the coinductive tree semantics [?], given by composition of the saturated semantics with desaturation. In fact it represents a different approach to the non-compositionality problem: instead of relaxing naturality to lax naturality, we simply forget about all the arrows of the index category \mathbf{L}_Σ^{op} , shifting the framework from $\mathbf{Set}^{\mathbf{L}_\Sigma^{op}}$ to $\mathbf{Set}^{|\mathbf{L}_\Sigma^{op}|}$. The substitutions on trees (that are essential, for instance, for the resolution algorithm given in [?]) exist at the saturated level, i.e. in $\mathcal{C}(\mathcal{S})(At)$, and they are given precisely as the operator $\bar{\theta}$ described at the end of Section ??.

Example 6. The coinductive tree for $\mathbf{List}(cons(x_1, x_2))$ in $\mathbf{NatList}$ is depicted on the right. It is constructed by desaturating the tree $\llbracket \mathbf{List}(cons(x_1, x_2)) \rrbracket_{p^\sharp}$ in Example ??, i.e., by pruning all the or-nodes (and their descendants) that are not labeled with id_2 .



5 Soundness and Completeness

The notion of coinductive tree leads to a semantics that is sound and complete with respect to SLD-resolution [?, Th.4.8]. To this aim, a key role is played by *derivation subtrees* of a given coinductive tree.

Definition 4. Let T be the n -coinductive tree for an atom A in a program \mathbb{P} . A subtree T' of T is a derivation subtree if it satisfies the following conditions:

1. the root of T' is the root of T ;
2. if an and-node of T belongs to T' , then just one of its children belongs to T' ;
3. if an or-node of T belongs to T' , then all its children belong to T' .

A refutation subtree (called *success subtree* in [?]) is a finite derivation subtree with only or-nodes as leaves.

In analogy with coinductive trees, we want to define a notion of subtree for saturated semantics. This requires care: saturated trees are associated with unification, which is more liberal than term-matching. In particular, similarly to and-or trees, they may represent unsound derivation strategies. However, in saturated trees *all* unifiers, and not just the most general ones, are taken into account. This gives enough flexibility to shape a sound notion of subtree, based on an implicit synchronization of the substitutions used in different branches.

Definition 5. Let T be the saturated tree for an atom A in a program \mathbb{P} . A subtree T' of T is called a synched derivation subtree if it satisfies properties 1-3 of Definition ?? and the following condition:

4. all or-nodes of T' at the same depth are labeled with the same substitution.

A synched refutation subtree is a finite synched derivation subtree with only or-nodes as leaves. Its answer is the substitution $\theta_{2k+1} \circ \dots \circ \theta_3 \circ \theta_1$, where θ_i is the (unique) substitution labeling the or-nodes of depth i and $2k + 1$ is its maximal depth.

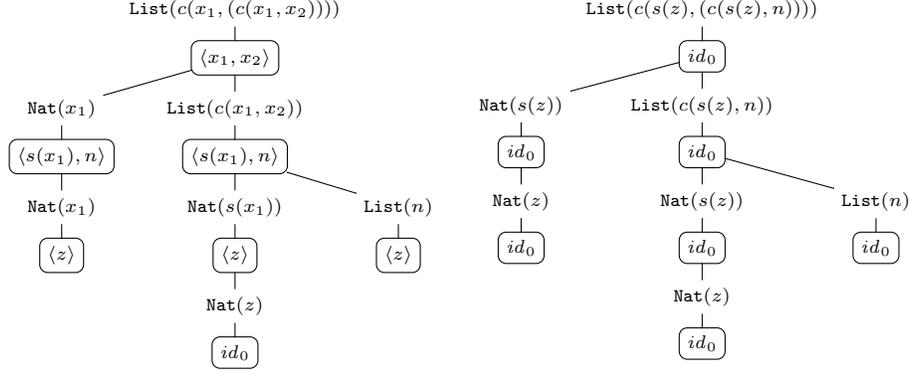


Fig. 1. Successful synched derivation subtrees for $\text{List}(\text{cons}(x_1, (\text{cons}(x_1, x_2))))$ (left) and $\text{List}(\text{cons}(\text{succ}(\text{zero}), (\text{cons}(\text{succ}(\text{zero}), \text{nil})))$ (right) in NatList . The symbols cons , nil , succ and zero are abbreviated to c , n , s and z respectively.

The prefix “synched” emphasizes the restriction to and-parallelism which is encoded in Definition ???. Intuitively, we force all subgoals at the same depth to proceed with *the same* substitution. For instance, this rules out the unsound derivation of [?, Ex.5.2].

Note that derivation subtrees can be seen as special instances of synched derivation subtrees where all the substitutions are forced to be identities.

Theorem 4 (Soundness and Completeness). *Let \mathbb{P} be a logic program and $A \in \text{At}(n)$ an atom. The following are equivalent.*

1. *The saturated tree for A in \mathbb{P} has a synched refutation subtree with answer θ .*
2. *There is some natural number m such that the m -coinductive tree for $A\theta$ in \mathbb{P} has a refutation subtree.*
3. *There is an SLD-refutation for $\{A\}$ in \mathbb{P} with computed answer τ such that there exists a substitution σ with $\sigma \circ \tau = \theta$.*

The statement $(2 \Leftrightarrow 3)$ is a rephrasing of [?, Th.4.8], while $(1 \Leftrightarrow 2)$ follows by compositionality and desaturation (Theorems ?? and ??).

Figure ?? provides an example of the argument for direction $(1 \Rightarrow 2)$. Note that the root of the rightmost tree is labeled with an atom of the form $A\theta$, where θ and A are respectively the answer and the label of the root of the leftmost tree. The key observation is that the rightmost tree is a refutation subtree of the 0-coinductive tree for $A\theta$ and can be obtained from the leftmost tree by a procedure involving the operator $\bar{\theta}$ discussed at the end of Section ??.

References

1. J. Adámek and V. Koubek. On the greatest fixed point of a set functor. *Theor. Comput. Sci.*, 150:57–75, 1995.

2. G. Amato, J. Lipton, and R. McGrail. On the algebraic structure of declarative programming languages. *Theor. Comput. Sci.*, 410(46):4626–4671, 2009.
3. F. Bonchi, M. G. Buscemi, V. Ciancia, and F. Gadducci. A presheaf environment for the explicit fusion calculus. *J. Autom. Reason.*, 49(2):161–183, 2012.
4. F. Bonchi and U. Montanari. Coalgebraic symbolic semantics. In *CALCO*, volume 5728 of *LNCS*, pages 173–190. Springer, 2009.
5. F. Bonchi and U. Montanari. Reactive systems, (semi-)saturated semantics and coalgebras on presheaves. *Theor. Comput. Sci.*, 410(41):4044–4066, 2009.
6. R. Bruni, U. Montanari, and F. Rossi. An interactive semantics of logic programming. *Theory and Practice of Logic Programming*, 1(6):647–690, 2001.
7. A. Corradini, M. Große-Rhode, and R. Heckel. A coalgebraic presentation of structured transition systems. *Theor. Comput. Sci.*, 260(1-2):27–55, 2001.
8. A. Corradini, R. Heckel, and U. Montanari. From SOS specifications to structured coalgebras: How to make bisimulation a congruence. *ENTCS*, 19(0):118 – 141, 1999.
9. A. Corradini and U. Montanari. An algebraic semantics for structured transition systems and its application to logic programs. *Theor. Comput. Sci.*, 103(1):51 – 106, 1992.
10. C. Dwork, P. C. Kanellakis, and J. C. Mitchell. On the sequential nature of unification. *The Journal of Logic Programming*, 1(1):35 – 50, 1984.
11. M. P. Fiore, E. Moggi, and D. Sangiorgi. A fully abstract model for the π -calculus. *Inf. Comput.*, 179(1):76–117, 2002.
12. M. P. Fiore and S. Staton. A congruence rule format for name-passing process calculi from mathematical structural operational semantics. In *LICS*, pages 49–58. IEEE, 2006.
13. J. A. Goguen. What is unification? - a categorical view of substitution, equation and solution. In *Resolution of Equations in Algebraic Structures, Volume 1: Algebraic Techniques*, pages 217–261. Academic, 1989.
14. G. Gupta and V. S. Costa. Optimal implementation of and-or parallel prolog. *Future Generation Computer Systems*, 10(1):71–92, 1994.
15. Y. Kinoshita and A. J. Power. A fibrational semantics for logic programs. In *Extensions of Logic Programming*, pages 177–191. Springer, 1996.
16. E. Komendantskaya, G. McCusker, and J. Power. Coalgebraic semantics for parallel derivation strategies in logic programming. In *AMAST*, volume 6486 of *LNCS*, pages 111–127. Springer, 2011.
17. E. Komendantskaya and J. Power. Coalgebraic derivations in logic programming. In *CSL*, volume 12 of *LIPICs*, pages 352–366. Schloss Dagstuhl, 2011.
18. E. Komendantskaya and J. Power. Coalgebraic semantics for derivations in logic programming. In *CALCO*, volume 6859 of *LNCS*, pages 268–282. Springer, 2011.
19. S. Mac Lane. *Categories for the Working Mathematician*. Graduate Texts in Mathematics. Springer, 2nd edition, Sept. 1998.
20. M. Miculan. A categorical model of the fusion calculus. *ENTCS*, 218:275–293, 2008.
21. M. Miculan and K. Yemane. A unifying model of variables and names. In *FOS-SACS*, volume 3441 of *LNCS*, pages 170–186. Springer, 2005.
22. U. Montanari and M. Sammartino. A network-conscious pi-calculus and its coalgebraic semantics. *Submitted to TCS (Festschrift for Glynn Winskel)*.
23. D. Sangiorgi. A theory of bisimulation for the pi-calculus. *Acta Inf.*, 33(1):69–97, 1996.
24. S. Staton. Relating coalgebraic notions of bisimulation. In *CALCO*, volume 5728 of *LNCS*, pages 191–205. Springer, 2009.
25. J. Worrell. Terminal sequences for accessible endofunctors. *ENTCS*, 19:24–38, 1999.