

Constructive Negation in Extensional Higher-Order Logic Programming

Angelos Charalambidis

Dept. of Informatics and Telecommunications
University of Athens
a.charalambidis@di.uoa.gr

Panos Rondogiannis

Dept. of Informatics and Telecommunications
University of Athens
prondo@di.uoa.gr

Abstract

Extensional higher-order logic programming has been recently proposed as an interesting extension of classical logic programming. An important characteristic of the new paradigm is that it preserves all the well-known properties of traditional logic programming. In this paper we enhance extensional higher-order logic programming with *constructive negation*. We argue that the main ideas underlying constructive negation are quite close to the existing proof procedure for extensional higher-order logic programming and for this reason the two notions amalgamate quite conveniently. We demonstrate the soundness of the resulting proof procedure and describe an actual implementation of a language that embodies the above ideas. In this way we obtain the first (to our knowledge) higher-order logic programming language supporting constructive negation and offering a new style of programming that genuinely extends that of traditional logic programming.

1 Introduction

Extensional higher-order logic programming has been recently proposed (Charalambidis et al. 2010; 2013) as an interesting extension of classical logic programming. The key idea behind the new paradigm is that all predicates defined in a program (even the higher-order ones) denote sets and therefore one can use standard extensional set theory in order to understand the meaning of programs and reason about them. A consequence of this fact is that the semantics and the proof theory of extensional higher-order logic programming smoothly extend the corresponding ones for traditional (ie., first-order) logic programming.

In this paper we extend the new paradigm of extensional higher-order logic programming with *constructive negation* (Chan 1988; 1989). The combination of higher-order characteristics and constructive negation in the new language, allows the programmer to specify in a compact way many interesting problems. For example, in the new language it is quite convenient to enumerate the sets that satisfy a specific property (such as all the colorings of a graph, all the subsets of a set that satisfy a condition, and so on). In order to theoretically justify the proposed higher-order language, we define its *completion semantics* and demonstrate

the soundness of the proof procedure with respect to this semantics. Finally, we describe an implementation of the new language which embodies the above ideas. In conclusion, the main contributions of the present work are the following:

- We propose the first (to our knowledge) higher-order logic programming language with constructive negation. The amalgamation of the two ideas seems quite natural and expressive (but is non-trivial from a technical point of view).
- We provide a theoretical justification of our proposal by defining the completion semantics for higher-order logic programs and by demonstrating the soundness of the proposed proof procedure with respect to this semantics. In this way we provide a promising direction for the study of the interplay between higher-order logic programming and negation-as-failure.

The rest of the paper is organized as follows. Section 2 describes the motivation behind the present work. Section 3 provides the syntax and semantics of the proposed higher-order language and Section 4 presents the corresponding proof-procedure and discusses its implementation. Section 5 demonstrates the soundness of the proof-procedure. Section 6 concludes the paper with discussion on future work.

2 Motivation

The purpose of this paper is to enhance extensional higher-order logic programming with constructive negation. These two notions prove to be very compatible, for reasons that we are going to explain in detail below.

In (Charalambidis et al. 2013) the higher-order logic programming language $\mathcal{H}$ is defined. Intuitively, $\mathcal{H}$ extends classical logic programming with higher-order user-defined predicates. The main novel characteristic of $\mathcal{H}$ with respect to other existing higher-order logic programming languages (such as Hilog (Chen, Kifer, and Warren 1993) and λ -Prolog (Miller and Nadathur 1986)) is that it has a purely extensional semantics: program predicates denote sets of objects and the semantics of the language is an extension of the traditional semantics of classical logic programming. Since predicates in $\mathcal{H}$ denote sets, the programmer is allowed to use uninstantiated predicate variables in clause bodies and in queries; when the proof-procedure of $\mathcal{H}$ encounters such

a variable, it performs a systematic enumeration of the possible values that this variable can take. For further details regarding $\mathcal{H}$ the interested reader should consult (Charalambidis et al. 2013). The following two examples motivate the language. To simplify our presentation for this section, we use an extended Prolog-like syntax. The exact syntax of the language that will be the object of our study will be specified in the next section.

Example 1. Consider the program:

$$p(Q) : -Q(0), Q(1).$$

The meaning of p can be understood in set-theoretic terms, i.e., p is a predicate that is true of all relations that contain (at least) 0 and 1. Thus, it is meaningful to ask the query:

$$?-p(R).$$

which will respond with an answer of the form $R = \{0, 1\} \cup L$ (having the obvious meaning described above).

Example 2. Consider the following program stating that a band is a group that has at least a singer and a guitarist:

$$\text{band}(B) : -\text{singer}(S), B(S), \text{guitarist}(G), B(G).$$

Suppose that we also have a database of musicians:

$$\begin{aligned} &\text{singer}(\text{sally}). \\ &\text{singer}(\text{steve}). \\ &\text{guitarist}(\text{george}). \\ &\text{guitarist}(\text{grace}). \end{aligned}$$

We can then ask the following query:

$$?- \text{band}(B).$$

which returns answers of the form $B = \{\text{sally}, \text{george}\} \cup L$, and so on.

It should be noted that in $\mathcal{H}$ one can ask queries with free predicate variables of any *order*; this means that we can have queries that return sets of sets of objects, and so on.

As demonstrated in (Charalambidis et al. 2013), one can define a sound and complete proof procedure for $\mathcal{H}$ and based on it to obtain a reasonably efficient implementation of the language. There are however certain aspects of the language that require further investigation. One basic property of all the higher-order predicates that can be defined in $\mathcal{H}$ is that they are *monotonic*. Intuitively, the monotonicity property states that if a predicate is true of a relation R then it is also true of every superset of R . In the above example, it is clear that if band is true of a relation B then it is also true of any band that is a superset of B . However, there are many natural higher-order predicates that are *non-monotonic*. For example, consider the predicate $\text{disconnected}(G)$ which succeeds if its input argument, namely a graph G , is disconnected. A graph is simply a set of pairs, and therefore disconnected is a second-order predicate. Notice now that disconnected is obviously non-monotonic: given graphs $G1$ and $G2$ with $G1 \subseteq G2$, it is possible that $\text{disconnected}(G1)$ succeeds but $\text{disconnected}(G2)$ fails.

The obvious idea in order to add non-monotonicity to $\mathcal{H}$ is to enhance the language with negation-as-failure. However, this is not as straightforward as it sounds, because even the simpler higher-order programs with negation face the well-known problem of *floundering* (Lloyd 1987). In classical logic programming a computation is said to *flounder* if at some point a goal is reached that contains only non-ground negative literals. In higher-order logic programming this problem is even more extended since a goal may additionally contain uninstantiated higher-order variables.

Example 3. Consider the following simple modification to the program of Example 1:

$$p(Q) : -Q(0), \text{not}(Q(1)).$$

Intuitively, p is true of all relations that contain 0 but they do not contain 1. The query:

$$?-p(R).$$

will not behave correctly if a simple-minded implementation of negation-as-failure is followed; this is due to the fact that in order to answer the given query one has to answer the subquery $\text{not}(R(1))$. Since the relation R is not fully known at the point of this call, it is not obvious how the implementation should treat such a subquery.

Constructive negation (Chan 1988; 1989) bypasses the problem of floundering. The main idea can be explained by a simple example. Consider the program:

$$\begin{aligned} &p(1). \\ &p(2). \end{aligned}$$

and the query:

$$?- \text{not}(p(X)).$$

The original idea (Chan 1988) behind constructive negation is that in order to answer a negative query that contains uninstantiated variables, the following procedure must be applied: we run the positive version of the query and we collect the solutions as a disjunction; we then return the negation of the disjunction as the answer to the original query. This idea can be extended even to the case where a query has an infinite number of answers (Chan 1989). In our example, the corresponding positive query (namely $?-p(X)$) has the answers $X=1$ and $X=2$. The negation of their disjunction is the conjunction $(X \neq 1) \wedge (X \neq 2)$. Observe now that the procedure behind constructive negation returns as answers not only substitutions (as it happens in negationless logic programming) but also inequalities.

Generalizing the above idea to the higher-order setting requires the ability to express some form of inequalities regarding the elements of sets. Intuitively, we would like to express the property that some element *does not belong to a set*. For example, given the query in Example 3, a reasonable answer would be $R = \{0\} \cup \{X \mid X \neq 1\}$.

We start with the syntax of $\mathcal{H}$ as introduced in (Charalambidis et al. 2013) and extend it in order to allow negation in clause bodies. The new language, $\mathcal{H}_{cn}$, has only one limitation with respect to the initial $\mathcal{H}$: the existential predicate variables that appear in queries or in clause bodies of

$\mathcal{H}_{cn}$ can only be sets of terms (and not, for example, sets of sets, or sets of sets of sets, and so on). The reasons for our restriction are two-fold: first, it is not straightforward to express inequalities regarding the elements of a set of sets; and second, it would be much more difficult to express the proof procedure if the queries were not restricted (and the implementation would also face significant complications).

It turns out that $\mathcal{H}_{cn}$ is quite appropriate in expressing problems that require a generate-and-test of sets in order to reach a solution. For example, many graph problems require the enumeration and testing of subgraphs of a given graph. In this respect, the language $\mathcal{H}_{cn}$ appears to have an application domain that is similar to that of modern ASP languages, which are being used for specifying computationally difficult problems (Gelfond and Lifschitz 1988; Lifschitz 2008). Of course, the higher-order approach is completely different than the ASP one in terms of syntax and semantics, and a comparative evaluation of the two paradigms deserves further investigation.

Example 4. Consider the following program which defines the subset relation:

```
subset(P, Q) :- not(non_subset(P, Q)).
non_subset(P, Q) :- P(X), not(Q(X)).
```

and assume that we have a predicate q which is true of 0, 1 and 2. Then, given the query

```
?-subset(P, q).
```

the interpreter produces all the subsets of the set $\{0, 1, 2\}$.

Example 5. Consider the following program which defines the relation $\text{twocolor}(G, R)$ which is true if R is a subset of the set of vertices of G that can be painted with the same color in a two-coloring of G .

```
twocolor(G, R) :- subset(R, vertex_set(G)),
                  not(non_twocolor(G, R)).
```

```
vertex_set(G)(X) :- G(X, _).
vertex_set(G)(X) :- G(_, X).
```

```
non_twocolor(G, R) :- G(X, Y), R(X), R(Y).
non_twocolor(G, R) :- G(X, Y), not(R(X)), not(R(Y)).
```

Assume that g is a binary predicate defining a graph. By asking the query:

```
?-twocolor(g, R).
```

the implementation will enumerate all the possible relations R that constitute valid two-colorings of g .

Example 6. The following program can be used to enumerate all the subgraphs of a given graph that are cliques:

```
clique(G, R) :- subset(R, vertex_set(G)),
                not(non_clique(G, R)).
```

```
non_clique(G, R) :- R(X), R(Y), not(G(X, Y)).
```

3 Syntax and Semantics

This section defines in a formal way the syntax and the semantics of the higher-order language $\mathcal{H}_{cn}$.

The Syntax of $\mathcal{H}_{cn}$

The language $\mathcal{H}_{cn}$ is in fact a more austere version of the ad-hoc language used for the examples of the previous section.

Example 7. The subset predicate of Example 4 can be expressed in $\mathcal{H}_{cn}$ as follows:

$$\text{subset} \leftarrow \lambda P. \lambda Q. \sim \exists X((P X) \wedge \sim(Q X))$$

The subset predicate is defined by a λ -expression (which obviates the need to have the formal parameters of the predicate in the left-hand side of the definition). Moreover, in the right-hand side we have an explicit existential quantifier for the variable X (in Prolog, if a variable appears in the body of a clause but not in the head, then it is implicitly existentially quantified). Consider the query:

```
← subset R p
```

Notice that here the variable R is not explicitly existentially quantified: we assume that the free variables of the goal are the ones for which the proof procedure will attempt to find bindings in order for the goal to be satisfied. If a variable appears explicitly existentially quantified in the goal, then the implementation will not produce a binding for it.

The type system of $\mathcal{H}_{cn}$ is based on two base types: o , the type of the boolean domain, and ι , the type of the domain of individuals (data objects). For example, every classical logic programming term is of type ι .

The basic difference in the types of $\mathcal{H}_{cn}$ from those of $\mathcal{H}$ is the existence of a type μ which restricts the set of predicate variables that can be existentially quantified or appear free in goal clauses (as explained in the previous section).

Definition 1. A type can either be *functional*, *argument*, or *predicate*, denoted by σ , ρ and π and defined as:

$$\begin{aligned} \sigma &:= \iota \mid (\iota \rightarrow \sigma) \\ \rho &:= \iota \mid \pi \\ \pi &:= o \mid (\rho \rightarrow \pi) \end{aligned}$$

The subtypes μ and κ of ρ and π are defined as follows:

$$\begin{aligned} \mu &:= \iota \mid \kappa \\ \kappa &:= \iota \rightarrow o \mid (\iota \rightarrow \kappa) \end{aligned}$$

The types μ and κ will be called *existential* and *set* types respectively.

The operator $\rightarrow$ is right-associative. A functional type that is different from ι will often be written in the form $\iota^n \rightarrow \iota$, $n \geq 1$. Similarly, a set type will often be written in the form $\iota^n \rightarrow o$, $n \geq 1$. It can be easily seen that every predicate type π can be written in the form $\rho_1 \rightarrow \dots \rightarrow \rho_n \rightarrow o$, $n \geq 0$ (for $n = 0$ we assume that $\pi = o$). We proceed with the alphabet of $\mathcal{H}_{cn}$:

Definition 2. The *alphabet* of $\mathcal{H}_{cn}$ consists of:

1. *Predicate variables* of every predicate type π (denoted by uppercase letters such as $R, Q, \dots$).
2. *Predicate constants* of every predicate type π (denoted by lowercase letters such as $p, q, r, \dots$).
3. *Individual variables* of type ι (denoted by uppercase letters such as $X, Y, Z, \dots$).

4. *Individual constants* of type ι (denoted by lowercase letters such as $a, b, c, \dots$).
5. *Function symbols* of every functional type $\sigma \neq \iota$ (denoted by lowercase letters such as $f, g, h, \dots$).
6. The following logical constant symbols: the propositional constants false and true of type o ; the equality constant $\approx$ of type $\iota \rightarrow \iota \rightarrow o$; the generalized disjunction and conjunction constants $\bigvee_\pi$ and $\bigwedge_\pi$ of type $\pi \rightarrow \pi \rightarrow \pi$, for every predicate type π ; the generalized inverse implication constants $\leftarrow_\pi$ of type $\pi \rightarrow \pi \rightarrow o$, for every predicate type π ; the equivalence constants $\leftrightarrow_\pi$ of type $\pi \rightarrow \pi \rightarrow o$, for every predicate type π ; the existential quantifiers $\exists_\mu$ of type $(\mu \rightarrow o) \rightarrow o$, for every existential type μ ; the negation constant $\sim$ of type $o \rightarrow o$.
7. The abstractor λ and the parentheses “(” and “)”.

The set consisting of the predicate variables and the individual variables of $\mathcal{H}_{cn}$ will be called the set of *argument variables* of $\mathcal{H}_{cn}$. Argument variables will be usually denoted by V, U and their subscripted versions.

Definition 3. The set of *body expressions* of the higher-order language $\mathcal{H}_{cn}$ is recursively defined as follows:

1. Every predicate variable (respectively, predicate constant) of type π is a body expression of type π ; every individual variable (respectively, individual constant) of type ι is a body expression of type ι ; the propositional constants false and true are body expressions of type o .
2. If f is an n -ary function symbol and $E_1, \dots, E_n$ are body expressions of type ι , then $(f E_1 \dots E_n)$ is a body expression of type ι .
3. If E_1 is a body expression of type $\rho \rightarrow \pi$ and E_2 is a body expression of type ρ , then $(E_1 E_2)$ is a body expression of type π .
4. If V is an argument variable of type ρ and E is a body expression of type π , then $(\lambda V.E)$ is a body expression of type $\rho \rightarrow \pi$.
5. If E_1, E_2 are expressions of type π , then $(E_1 \bigwedge_\pi E_2)$ and $(E_1 \bigvee_\pi E_2)$ are body expressions of type π .
6. If E_1, E_2 are expressions of type ι , then $(E_1 \approx E_2)$ is a body expression of type o .
7. If E is an expression of type o and V is an existential variable of type μ , then $(\exists_\mu V E)$ is a body expression of type o .
8. If E is a body expression of type o , then $(\sim E)$ is a body expression of type o .

The type subscripts (such as μ in $\exists_\mu$) will often be omitted when they are obvious or immaterial in the foregoing discussion. Moreover, we will write $\wedge$ and $\vee$ instead of $\bigwedge_o$ and $\bigvee_o$. We will use *type*(E) to denote the type of the body expression E . The notions of *free* and *bound* variables of a body expression are defined as usual. A body expression is called *closed* if it does not contain any free variables. Given an expression E , we denote by $fv(E)$ the set of all free variables of E . By overloading notation, we will also write $fv(S)$, where S is a set of expressions.

We will often write $\hat{A}$ to denote a (possibly empty) sequence $\langle A_1, \dots, A_n \rangle$ of expressions. For example we

will write $(E \hat{A})$ to denote an application $(E A_1 \dots A_n)$; $(\lambda \hat{X}.E)$ to denote $(\lambda X_1. \dots \lambda X_n.E)$; $(\exists \hat{V} E)$ to denote $(\exists V_1 \dots \exists V_n E)$. When $\hat{A}$ or $\hat{X}$ or $\hat{V}$ is empty, then the corresponding expression is considered identical to E .

A body expression of the form $(E_1 \approx E_2)$ will be called an *equality* while an expression of the form $\sim \exists \hat{V}(E_1 \approx E_2)$ will be called an *inequality*; in the latter case $\hat{V}$ may be empty (in which case the inequality is of the form $\sim(E_1 \approx E_2)$). Let $\hat{E} = \langle E_1, \dots, E_n \rangle$ and $\hat{E}' = \langle E'_1, \dots, E'_n \rangle$, where all E_i, E'_i are of type ι , we will write $(\hat{E} \approx \hat{E}')$ to denote the expression $(E_1 \approx E'_1) \wedge \dots \wedge (E_n \approx E'_n)$; if $n = 0$, then the conjunction is the constant true.

Definition 4. The set of *clausal expressions* of the higher-order language $\mathcal{H}_{cn}$ is defined as follows:

1. If p is a predicate constant of type π and E is a closed body expression of type π then $p \leftarrow_\pi E$ is a clausal expression of $\mathcal{H}_{cn}$, also called a *program clause*.
2. If E is a body expression of type o and each free variable in E is of type μ then false $\leftarrow_o E$ (usually denoted by $\leftarrow_o E$ or just $\leftarrow E$) is a clausal expression of $\mathcal{H}_{cn}$, also called a *goal clause*.
3. If p is a predicate constant of type π and E is a closed body expression of type π then $p \leftrightarrow_\pi E$ is a clausal expression of $\mathcal{H}_{cn}$, also called a *completion expression*.

All clausal expressions of $\mathcal{H}_{cn}$ have type o .

Notice that (following the tradition of first-order logic programming) we will talk about the “empty clause” which is denoted by $\square$ and is equivalent to $\leftarrow$ true.

Definition 5. A program of $\mathcal{H}_{cn}$ is a finite set of program clauses of $\mathcal{H}_{cn}$.

We proceed by defining the completion of a program:

Definition 6. Let P be a program and let p be a predicate constant of type π . Then, the *completed definition* for p with respect to P is obtained as follows:

- if there exist exactly $k > 0$ program clauses of the form $p \leftarrow_\pi E_i$ where $i \in \{1, \dots, k\}$ for p in P , then the completed definition for p is the expression $p \leftrightarrow_\pi E$, where $E = E_1 \bigvee_\pi \dots \bigvee_\pi E_k$.
- if there are no program clauses for p in P then the completed definition for p is the expression $p \leftrightarrow_\pi E$, where E is of type π and $E = \lambda \hat{X}. \text{false}$.

The expression E on the right-hand side of the completed definition of p will be called the *completed expression* for p with respect to P .

We can now extend the notion of completion (Clark 1977; Lloyd 1987) to apply to our higher-order programs:

Definition 7. Let P be a program. Then, the *completion comp*(P) of P is the set consisting of all the completed definitions for all predicate constants that appear in P .

The Semantics of $\mathcal{H}_{cn}$

In this section we introduce the semantics of $\mathcal{H}_{cn}$. We start by defining the meaning of types of $\mathcal{H}_{cn}$; for the predicate types we also define a corresponding partial order (which will be used in the definition of the semantics of $\mathcal{H}_{cn}$).

Definition 8. Let D be a non-empty set. Then the meanings of the types of our language can be specified as follows:

- $\llbracket \iota \rrbracket_D = D$;
- $\llbracket \iota^n \rightarrow \iota \rrbracket_D = D^n \rightarrow D$;
- $\llbracket o \rrbracket_D = \{false, true\}$. A partial order $\sqsubseteq_o$ can be defined on the set $\llbracket o \rrbracket_D$ as follows: $false \sqsubseteq_o true$, $false \sqsubseteq_o false$ and $true \sqsubseteq_o true$;
- $\llbracket \rho \rightarrow \pi \rrbracket_D = \llbracket \rho \rrbracket_D \rightarrow \llbracket \pi \rrbracket_D$. A partial order $\sqsubseteq_{\rho \rightarrow \pi}$ can be defined as follows: for all $f, g \in \llbracket \rho \rightarrow \pi \rrbracket_D$, $f \sqsubseteq_{\rho \rightarrow \pi} g$ if and only if $f(d) \sqsubseteq_{\pi} g(d)$, for all $d \in \llbracket \rho \rrbracket_D$.

It can be easily shown that for every predicate type π , $\llbracket \pi \rrbracket_D$ is a complete lattice, and therefore for every subset of $\llbracket \pi \rrbracket_D$ there exists a *lub* function (denoted by $\sqcup_{\pi}$) and a *glb* function (denoted by $\sqcap_{\pi}$).

The notions of *interpretation* and *state* can be defined in a standard way:

Definition 9. An interpretation I of $\mathcal{H}_{cn}$ consists of:

1. a nonempty set D , called the *domain* of I
2. an assignment to each individual constant symbol c , of an element $I(c) \in D$
3. an assignment to each predicate constant p of type π , of an element $I(p) \in \llbracket \pi \rrbracket_D$
4. an assignment to each function symbol f of type $\iota^n \rightarrow \iota$, of a function $I(f) \in D^n \rightarrow D$.

Definition 10. Let D be a nonempty set. Then, a *state* s of $\mathcal{H}_{cn}$ over D is a function that assigns to each argument variable V of type ρ of $\mathcal{H}_{cn}$ an element $s(V) \in \llbracket \rho \rrbracket_D$.

In the following, $s[V/d]$ is used to denote a state that is identical to s the only difference being that the new state assigns to V the value d .

Definition 11. Let I be an interpretation of $\mathcal{H}_{cn}$, let D be the domain of I and let s be a state over D . Then, the semantics of expressions of $\mathcal{H}_{cn}$ with respect to I and s is defined as follows:

1. $\llbracket false \rrbracket_s(I) = false$
2. $\llbracket true \rrbracket_s(I) = true$
3. $\llbracket c \rrbracket_s(I) = I(c)$, for every individual constant c
4. $\llbracket p \rrbracket_s(I) = I(p)$, for every predicate constant p
5. $\llbracket V \rrbracket_s(I) = s(V)$, for every argument variable V
6. $\llbracket (f E_1 \dots E_n) \rrbracket_s(I) = I(f) \llbracket E_1 \rrbracket_s(I) \dots \llbracket E_n \rrbracket_s(I)$, for every n -ary function symbol f
7. $\llbracket (E_1 E_2) \rrbracket_s(I) = \llbracket E_1 \rrbracket_s(I) (\llbracket E_2 \rrbracket_s(I))$
8. $\llbracket (\lambda V. E) \rrbracket_s(I) = \lambda d. \llbracket E \rrbracket_{s[V/d]}(I)$, where d ranges over $\llbracket type(V) \rrbracket_D$
9. $\llbracket (E_1 \vee_{\pi} E_2) \rrbracket_s(I) = \sqcup_{\pi} \{ \llbracket E_1 \rrbracket_s(I), \llbracket E_2 \rrbracket_s(I) \}$, where $\sqcup_{\pi}$ is the least upper bound function on $\llbracket \pi \rrbracket_D$
10. $\llbracket (E_1 \wedge_{\pi} E_2) \rrbracket_s(I) = \sqcap_{\pi} \{ \llbracket E_1 \rrbracket_s(I), \llbracket E_2 \rrbracket_s(I) \}$, where $\sqcap_{\pi}$ is the greatest lower bound function on $\llbracket \pi \rrbracket_D$
11. $\llbracket (E_1 \approx E_2) \rrbracket_s(I) = \begin{cases} true, & \text{if } \llbracket E_1 \rrbracket_s(I) = \llbracket E_2 \rrbracket_s(I) \\ false, & \text{otherwise} \end{cases}$
12. $\llbracket (\exists V E) \rrbracket_s(I) = \begin{cases} true, & \text{if there exists } d \in \llbracket type(V) \rrbracket_D \\ & \text{such that } \llbracket E \rrbracket_{s[V/d]}(I) = true \\ false, & \text{otherwise} \end{cases}$

13. $\llbracket (\sim E) \rrbracket_s(I) = \begin{cases} true, & \text{if } \llbracket E \rrbracket_s(I) = false \\ false, & \text{if } \llbracket E \rrbracket_s(I) = true \end{cases}$
14. $\llbracket p \leftarrow_{\pi} E \rrbracket_s(I) = \begin{cases} true, & \text{if } \llbracket E \rrbracket_s(I) \sqsubseteq_{\pi} I(p) \\ false, & \text{otherwise} \end{cases}$
15. $\llbracket p \leftrightarrow_{\pi} E \rrbracket_s(I) = \begin{cases} true, & \text{if } \llbracket E \rrbracket_s(I) = I(p) \\ false, & \text{otherwise} \end{cases}$
16. $\llbracket \leftarrow E \rrbracket_s(I) = \begin{cases} true, & \text{if } \llbracket E \rrbracket_s(I) = false \\ false, & \text{otherwise} \end{cases}$

Definition 12. Let S be a set of expressions of type o and let I be an interpretation of $\mathcal{H}_{cn}$. We say that I is a *model* of S if for every $E \in S$ and for every state s over the domain of I , $\llbracket E \rrbracket_s(I) = true$.

Clark's Equality Theory can also be added to our fragment as follows:

Definition 13. Let I be an interpretation of $\mathcal{H}_{cn}$. We will say that I satisfies *Clark's Equality Theory (CET)* if for all states s over the domain of I , all the following hold:

1. $\llbracket (f \hat{E}) \approx (f \hat{E}') \rrbracket_s(I) = \llbracket \hat{E} \approx \hat{E}' \rrbracket_s(I)$ for every n -ary function symbol f and for all $\hat{E} = \langle E_1, \dots, E_n \rangle$ and for all $\hat{E}' = \langle E'_1, \dots, E'_n \rangle$
2. $\llbracket c \approx d \rrbracket_s(I) = false$ for all different individual constants c and d
3. $\llbracket (f \hat{E}) \approx c \rrbracket_s(I) = false$ for every function symbol f and every individual constant c
4. $\llbracket (f \hat{E}) \approx (g \hat{E}') \rrbracket_s(I) = false$ for all different function symbols f and g and all $\hat{E}$ and $\hat{E}'$
5. $\llbracket X \approx E \rrbracket_s(I) = false$ if $X \in fv(E)$ and $X \neq E$

In the rest of this paper we will often talk about “models of the completion $comp(P)$ of a program P ”; in every such case we will implicitly assume that our models also satisfy Clark's Equality Theory.

4 Proof Procedure

In this section we define a proof procedure for $\mathcal{H}_{cn}$ and discuss its implementation. As in the case of classical logic programming, substitutions and unifiers also play an important role in our setting. However, they need to be significantly extended, as discussed below.

Substitutions and Unifiers

In classical logic programming, a substitution is a mapping from variables to terms. Since in our case queries may contain uninstantiated variables of set types, the notion of substitution must be extended to assign to these variables expressions that represent sets. In the following we define the notion of *basic expression* which generalizes the notion of term from classical logic programming. Intuitively, a basic element is either a term or a set of terms.

Definition 14. The set of *basic expressions* of $\mathcal{H}_{cn}$ of type μ is defined recursively as follows:

1. Every expression of $\mathcal{H}_{cn}$ of type ι is a basic expression of type ι .

2. Every predicate variable of type κ is a basic expression of type κ .
3. If E_1, E_2 are basic expressions of type κ then $E_1 \vee_{\kappa} E_2$ and $E_1 \wedge_{\kappa} E_2$ are basic expressions of type κ .
4. The expressions of the following form, are basic expressions of type $\iota^n \rightarrow o$:

- $\lambda\hat{X}. \exists\hat{V}(\hat{X} \approx \hat{A})$
- $\lambda\hat{X}. \sim\exists\hat{V}(\hat{X} \approx \hat{A})$

where $\hat{X} = \langle X_1, \dots, X_n \rangle$, $\hat{A} = \langle A_1, \dots, A_n \rangle$, each X_i is a variable of type ι , each A_i is a basic expressions of type ι and $\hat{V}$ is a possibly empty subset of $fv(\hat{A})$.

Example 8. The basic element $\lambda X. \sim\exists Y(X \approx Y)$ represents the empty set while $\lambda X. (X \approx a) \vee \lambda X. (X \approx (f b))$ represents the set $\{a, (f b)\}$. The basic element $\lambda X. \exists Y(X \approx (f Y))$ is the set that contains all the terms of the form $(f \dots)$. The basic element $\lambda X. \sim\exists Y(X \approx (f Y))$ is the set that contains all the terms that are not of the form $(f \dots)$. The basic element $\lambda X. \exists Y(X \approx (f Y)) \wedge \lambda X. \sim(X \approx (f b))$ is the set that contains all the terms of the form $(f \dots)$ with the exception of $(f b)$.

The definition of a substitution can be extended:

Definition 15. A substitution θ is a finite set of the form $\{V_1/E_1, \dots, V_n/E_n\}$ where V_i are different argument variables of $\mathcal{H}_{cn}$ and each E_i is a body expression of $\mathcal{H}_{cn}$ having the same type as V_i . We write $dom(\theta) = \{V_1, \dots, V_n\}$ and $range(\theta) = \{E_1, \dots, E_n\}$. A substitution is called *basic* if all E_i are basic expressions. A substitution is called *zero-order* if the type of V_i is ι for all $i \in \{1, \dots, n\}$. The substitution corresponding to the empty set will be called the *identity substitution* and will be denoted by ϵ .

In the following definition (and also in the rest of the paper) we assume that if expressions $E_1, \dots, E_n$ occur in a certain mathematical context, then in these expressions all bound variables are chosen to be different from the free variables. This is called the *the bound variable convention* in (Barendregt 1984, pages 26–27) (and is actually equivalent to the α -renaming operation of λ -calculus).

Definition 16. Let θ be a substitution and let E be a body expression. Then, $E\theta$ is an expression obtained as follows:

- $E\theta = E$, if E is false, true, c, or p
- $\forall\theta = \theta(V)$ if $V \in dom(\theta)$; otherwise, $\forall\theta = V$
- $(f E_1 \dots E_n)\theta = (f E_1\theta \dots E_n\theta)$
- $(E_1 E_2)\theta = (E_1\theta E_2\theta)$
- $(\lambda V.E)\theta = (\lambda V.E\theta)$
- $(E_1 \vee_{\pi} E_2)\theta = (E_1\theta \vee_{\pi} E_2\theta)$
- $(E_1 \wedge_{\pi} E_2)\theta = (E_1\theta \wedge_{\pi} E_2\theta)$
- $(E_1 \approx E_2)\theta = (E_1\theta \approx E_2\theta)$
- $(\exists V E)\theta = (\exists V E\theta)$
- $(\sim E)\theta = \sim E\theta$

In the following, if $\hat{E} = \langle E_1, \dots, E_n \rangle$, we will write $\hat{E}\theta$ to denote the sequence $\langle E_1\theta, \dots, E_n\theta \rangle$.

Definition 17. Let $\theta = \{V_1/E_1, \dots, V_m/E_m\}$ and $\sigma = \{V'_1/E'_1, \dots, V'_m/E'_m\}$ be substitutions. Then the composition $\theta\sigma$ of θ and σ is the substitution obtained from the set $\{V_1/E_1\sigma, \dots, V_m/E_m\sigma, V'_1/E'_1, \dots, V'_n/E'_n\}$ by deleting any $V_i/E_i\sigma$ for which $V_i = E_i\sigma$ and deleting any V'_j/E'_j for which $V'_j \in dom(\theta)$.

The notions of “unifier” and “most general unifier” are the same as in classical first-order logic programming.

Definition 18. Let S be a set of terms of $\mathcal{H}_{cn}$ (ie., expressions of type ι). A zero-order substitution θ will be called a *unifier* of the expressions in S if the set $S\theta = \{E\theta \mid E \in S\}$ is a singleton. The zero-order substitution θ will be called a *most general unifier* of S (denoted by $mgu(S)$), if for every unifier σ of the expressions in S , there exists a zero-order substitution γ such that $\sigma = \theta\gamma$.

The most general unifier of a set of terms S can be computed using the unification algorithm described in (Lloyd 1987). The output of the unification algorithm is denoted by $unify(S)$.

The following *substitution lemma* can easily be established by structural induction on E :

Lemma 1 (Substitution Lemma). *Let I be an interpretation of $\mathcal{H}_{cn}$ and let s be a state over the domain of I . Let θ be a basic substitution and E be a body expression. Then, $\llbracket E\theta \rrbracket_s(I) = \llbracket E \rrbracket_{s'}(I)$, where $s'(V) = \llbracket \theta(V) \rrbracket_s(I)$ if $V \in dom(\theta)$ and $s'(V) = s(V)$, otherwise.*

The above lemma will be used in the proof of the soundness of the resolution proof procedure that follows.

The Procedure

In this subsection we describe the actual proof procedure for $\mathcal{H}_{cn}$. Intuitively, we could say that our approach generalizes to the higher-order case the constructive negation procedure proposed by D. Chan in (Chan 1989). In order to define the overall proof procedure we utilize three definitions, described as follows:

- The central definition of our procedure is Definition 20 which defines the *single-step derivation* of goals. Intuitively, the single-step derivation is an extension of the single-step derivation of classical logic programming (see for example (Lloyd 1987, pages 50–41)) in order to apply to higher-order logic programming and to constructive negation. Notice that the single-step derivation is the only one among the three that may produce substitutions (whose final composition will lead to the computed answer for the initial goal).
- The *reduction* (Definition 21) is the simplest among the three definitions. Roughly speaking it corresponds to certain simple computational steps that advance the evaluation of the goal. For example, the β -reduction step for a λ -expression belongs to this reduction relation. Another simple reduction step is the replacement of a predicate constant with its defining expression in the program.
- The *negative reduction* (Definition 22) is the most complicated among the three. Intuitively, it corresponds to what D. Chan in (Chan 1989) refers to as “the negation of

answers” and the “normalization of answers”. The definition is more complicated in our case since it applies to the more demanding setting of higher-order logic programming.

One difference between the proof procedures for constructive negation (such as for example the one given in (Chan 1989)) and the traditional SLD-resolution for negationless logic programming, is the following: in classical SLD-resolution, a derivation is successful if we reach the empty goal; however, in constructive negation, a derivation may be successful if we reach a goal that consists of inequalities that can not be further simplified. Therefore, the termination condition for constructive negation is different than that of SLD-resolution. Moreover, an *answer* in classical resolution is simply a substitution for some of the variables of the goal. In constructive negation, an answer is a pair consisting of a substitution and of a conjunction of inequalities (intuitively, the inequalities that remain at the end of the derivation)¹. In the higher-order case that we are examining, an answer is also a pair, the only difference being that the domain of the substitution may contain higher-order variables.

Before we present the actual definitions, we need the following categorization of inequalities, which is pretty standard in the area of constructive negation (see for example (Chan 1989)):

Definition 19. An inequality $\sim\exists\hat{V}(E_1 \approx E_2)$ is considered

- *valid* if E_1 and E_2 cannot be unified;
- *unsatisfiable* if there is a substitution θ that unifies E_1 and E_2 and contains only bindings of variables in $\hat{V}$;
- *satisfiable* if it is not unsatisfiable.

An inequality will be called *primitive* if it is satisfiable, non valid and either E_1 or E_2 is a variable.

In the following definition we will refer to a conjunction $A_1 \wedge \dots \wedge A_n$ of expressions. We will assume that each A_i is not a conjunction itself. The conjunction of a single expression is the expression itself and the conjunction of zero expressions is the expression true.

Definition 20. Let P a program and let G_k and G_{k+1} be goal clauses. Moreover, let G_k be a conjunction $\leftarrow A_1 \wedge \dots \wedge A_n$, where each A_i is a body expression of type o . Let A_i be one of the $A_1, \dots, A_n$ and let us call it the *selected expression*. Let $A' = A_1 \wedge \dots \wedge A_{i-1} \wedge A_{i+1} \wedge \dots \wedge A_n$. We say that G_{k+1} is *derived in one step* from G_k using θ (denoted as $G_k \xrightarrow{\theta} G_{k+1}$) if one of the following conditions applies:

1. if A_i is true and $n > 1$ then $G_{k+1} = \leftarrow A'$ is derived from G_k using $\theta = \epsilon$;
2. if A_i is $(E_1 \vee E_2)$ then $G_{k+1} = \leftarrow A_1 \wedge \dots \wedge E_j \wedge \dots \wedge A_n$ is derived from G_k using $\theta = \epsilon$ where $j \in \{1, 2\}$;
3. if A_i is $(\exists V E)$ then $G_{k+1} = \leftarrow A_1 \wedge \dots \wedge E \wedge \dots \wedge A_n$ is derived² from G_k using $\theta = \epsilon$;

¹Actually, in (Chan 1989) the substitution is combined with the inequalities so as to obtain a sequence consisting of equalities and inequalities.

²Recall that by the Bound Variable Convention, V does not appear free in A' .

4. if $A_i \rightsquigarrow A'_i$ then $G_{k+1} = \leftarrow A_1 \wedge \dots \wedge A'_i \wedge \dots \wedge A_n$ is derived from G_k using $\theta = \epsilon$;
5. if A_i is $(E_1 \approx E_2)$ then $G_{k+1} = \leftarrow A'\theta$ is derived from G_k using $\theta = mgu(E_1, E_2)$;
6. if A_i is $(R \hat{E})$ and R is a predicate variable of type κ then $G_{k+1} = \leftarrow A'\theta$ is derived from G_k using $\theta = \{R/(\lambda\hat{X}.(\hat{X} \approx \hat{E}) \vee_{\kappa} R')\}$ where R' is a fresh³ predicate variable of type κ ;
7. if A_i is $\sim\exists\hat{V} E$ and A_i is negatively-reduced to A'_i then $G_{k+1} = \leftarrow A_1 \wedge \dots \wedge A'_i \wedge \dots \wedge A_n$ is derived from G_k using $\theta = \epsilon$;
8. if A_i is $\sim\exists\hat{V}(R \hat{E})$ and R is a predicate variable of type κ and $R \notin \hat{V}$ then $G_{k+1} = \leftarrow A'\theta$ is derived from G_k using $\theta = \{R/(\lambda\hat{X}. \sim\exists\hat{V}(\hat{X} \approx \hat{E}) \wedge_{\kappa} R')\}$ where R' is a fresh predicate variable of type κ ;
9. if A_i is $\sim\exists\hat{V} \sim(R \hat{E})$ and R is a predicate variable of type κ and $R \notin \hat{V}$ then $G_{k+1} = \leftarrow A'\theta$ is derived from G_k using $\theta = \{R/(\lambda\hat{X}. \exists\hat{V}(\hat{X} \approx \hat{E}) \vee_{\kappa} R')\}$ where R' is a fresh predicate variable of type κ .

Definition 21. Let P be a program and E, E' be body expressions of type o . We say that E is reduced (wrt. to P) to E' (denoted as $E \rightsquigarrow E'$) if one of the following conditions applies:

1. $p \hat{A} \rightsquigarrow E \hat{A}$, where E is the completed expression for p with respect to P ;
2. $(\lambda X.E) B \hat{A} \rightsquigarrow E\{X/B\} \hat{A}$;
3. $(E_1 \vee_{\pi} E_2) \hat{A} \rightsquigarrow (E_1 \hat{A}) \vee_o (E_2 \hat{A})$;
4. $(E_1 \wedge_{\pi} E_2) \hat{A} \rightsquigarrow (E_1 \hat{A}) \wedge_o (E_2 \hat{A})$.

Definition 22. Let P be a program and B, B' be body expressions where $B = \sim\exists\hat{U}(A_1 \wedge \dots \wedge A_n)$ and each A_i is a body expression except from conjunction. Let A_i be the selected expression and $A' = A_1 \wedge \dots \wedge A_{i-1} \wedge A_{i+1} \wedge \dots \wedge A_n$. Then, we say that B is *negatively-reduced* to B' if one of the following conditions applies:

1. if A_i is false then $B' = \text{true}$;
2. if A_i is true and $n = 1$ then $B' = \text{false}$ else $B' = \sim\exists\hat{U} A'$;
3. if A_i is $(E_1 \vee E_2)$ then $B' = B'_1 \wedge B'_2$ where $B'_j = \sim\exists\hat{U}(A_1 \wedge \dots \wedge E_j \wedge \dots \wedge A_n)$, $j \in \{1, 2\}$;
4. if A_i is $(\exists V E)$ then $B' = \sim\exists\hat{U} V(A_1 \wedge \dots \wedge E \wedge \dots \wedge A_n)$;⁴
5. if $A_i \rightsquigarrow A'_i$ then $B' = \sim\exists\hat{U}(A_1 \wedge \dots \wedge A'_i \wedge \dots \wedge A_n)$;
6. if A_i is $(E_1 \approx E_2)$, then
 - (a) if $\sim\exists\hat{U}(E_1 \approx E_2)$ is valid then $B' = \text{true}$;
 - (b) if $\sim\exists\hat{U}(E_1 \approx E_2)$ is non-valid and neither E_1 nor E_2 is a variable then $B' = \sim\exists\hat{U}(A_1 \wedge \dots \wedge A_{i-1} \wedge (\hat{X} \approx \hat{X}\theta) \wedge A_{i+1} \wedge \dots \wedge A_n)$ where $\theta = \text{unify}(E_1, E_2)$ and $\hat{X} = \text{dom}(\theta)$;

³This means a variable that has not appeared in goals and substitutions until this point of the derivation.

⁴Recall that by the Bound Variable Convention, V does not appear free in A' .

- (c) if $\sim \exists \hat{U}(E_1 \approx E_2)$ is unsatisfiable and either E_1 or E_2 is a variable in $\hat{U}$, then $B' = \sim \exists \hat{U}(A'\theta)$ where $\theta = \{X/E\}$ and X is the one expression that is a variable in $\hat{U}$ and E is the other;
- (d) if $\sim \exists \hat{U}(E_1 \approx E_2)$ is primitive and $n > 1$ then $B' = \sim \exists \hat{U}_1 A_i \vee \exists \hat{U}_1 (A_i \wedge \sim \exists \hat{U}_2 A')$ where $\hat{U}_1$ are the variables in $\hat{U}$ that are free in A_i and $\hat{U}_2$ the variables in $\hat{U}$ not in $\hat{U}_1$;
- 7. if A_i is $(R \hat{E})$ and R is a predicate variable, then
 - (a) if $R \in \hat{U}$ then $B' = \sim \exists \hat{U}'(A'\theta)$ where substitution $\theta = \{R/(\lambda X.(X \approx E) \vee_{\kappa} R')\}$, R' is a predicate variable of the same type with R and $\hat{U}'$ is the same with $\hat{U}$ but the variable R has been replaced with R' ;
 - (b) if $R \notin \hat{U}$ and $n > 1$ then $B' = \sim \exists \hat{U}_1 A_i \vee \exists \hat{U}_1 (A_i \wedge \sim \exists \hat{U}_2 A') \wedge B$ where $\hat{U}_1$ are the variables in $\hat{U}$ that are free in A_i and $\hat{U}_2$ the variables in $\hat{U}$ not in $\hat{U}_1$;
- 8. if A_i is $\sim \exists \hat{V} E$ and A_i is negatively-reduced to A'_i then $B' = \sim \exists \hat{U}(A_1 \wedge \dots \wedge A'_i \wedge \dots \wedge A_n)$;
- 9. if A_i is a primitive inequality $\sim \exists \hat{V}(E_1 \approx E_2)$ then
 - (a) if A_i contains free variables in $\hat{U}$ and A' is a conjunction of primitive inequalities then $B' = \sim \exists \hat{U} A'$;
 - (b) if A_i does not contain any free variables in $\hat{U}$ then $B' = \exists \hat{V}(E_1 \approx E_2) \vee \sim \exists \hat{U} A'$;
- 10. if A_i is $\sim \exists \hat{V}(R \hat{E})$ and $R \notin \hat{V}$ is a predicate variable, then
 - (a) if $R \in \hat{U}$ then $B' = \sim \exists \hat{U}'(A'\theta)$ where substitution $\theta = \{R/(\lambda X. \sim \exists \hat{V}(X \approx E) \wedge_{\kappa} R')\}$, R' is a predicate variable of the same type κ with R and $\hat{U}'$ is the same with $\hat{U}$ but the variable R has been replaced with R' ;
 - (b) if $R \notin \hat{U}$ and $n > 1$ then $B' = \sim \exists \hat{U}_1 A_i \vee \exists \hat{U}_1 (A_i \wedge \sim \exists \hat{U}_2 A') \wedge B$ where $\hat{U}_1$ are the variables in $\hat{U}$ that are free in A_i and $\hat{U}_2$ the variables in $\hat{U}$ not in $\hat{U}_1$;
 - (c) if $R \notin \hat{U}$, $n = 1$ and $\hat{V}$ is nonempty then $B' = \exists \hat{V} \sim \exists \hat{U}(\sim (R \hat{E}) \wedge \sim \exists \hat{V}'(R E'))$;
- 11. if A_i is $\sim \exists \hat{V} \sim (R \hat{E})$ and R is a predicate variable and $R \notin \hat{V}$, then
 - (a) if $R \in \hat{U}$ then $B' = \sim \exists \hat{U}'(A'\theta)$ where substitution $\theta = \{R/(\lambda X. \exists \hat{V}(X \approx E) \vee_{\kappa} R')\}$, R' is a predicate variable of the same type κ with R and $\hat{U}'$ is the same with $\hat{U}$ but the variable R has been replaced with R' ;
 - (b) if $R \notin \hat{U}$ and $n > 1$ then $B' = \sim \exists \hat{U}_1 A_i \vee \exists \hat{U}_1 (A_i \wedge \sim \exists \hat{U}_2 A') \wedge B$;
 - (c) if $R \notin \hat{U}$, $n = 1$ and $\hat{V}$ is nonempty then $B' = \exists \hat{V} \sim \exists \hat{U}((R \hat{E}) \wedge \sim \exists \hat{V}' \sim (R E'))$.

One can easily check that all the three above definitions preserve the main property of $\mathcal{H}_{cn}$ goals, namely that all the free variables in them have existential types.

We proceed by defining the notion of a *multi-step derivation* and by characterizing when a multi-step derivation is successful and the form of the computed answer. We first define the notions of terminal and primitive goals:

Definition 23. Let P be a program and $G \leftarrow A$ be a goal clause. Then, G will be called *terminal* if no goals can be derived in one step from G . The goal G is called *primitive* if A is a (possibly empty) conjunction of primitive inequalities.

Notice that every primitive goal is also terminal. Notice also that the empty goal $\square$ is primitive.

Definition 24. Let $G \leftarrow A_1 \wedge \dots \wedge A_n$ be a primitive goal and S be a set of variables. We say that G_S is a primitive goal obtained by restricting G to S when for each inequality A_i in G_S every free variable of A_i is in S .

Definition 25. Let P be a program and G be a goal. Assume that $P \cup \{G\}$ has a finite derivation $G_0 = G, G_1, \dots, G_n$ with basic substitutions $\theta_1, \dots, \theta_n$, such that the goal G_n is primitive. Then we will say that $P \cup \{G\}$ has a successful derivation of length n with primitive goal G_n and using basic substitution $\theta = \theta_1 \dots \theta_n$.

Notice that it can easily be verified that the composition of basic substitutions is indeed a basic substitution as implied by the above definition.

We can now define the important notion of *computed answer*:

Definition 26. Let P be a program, G be a goal and assume that $P \cup \{G\}$ has a successful derivation with primitive goal G' and basic substitution θ . Then (σ, G'') is a *computed answer* for $P \cup \{G\}$ where σ is the basic substitution obtained by restricting θ to the free variables of G and G'' is the primitive goal G' restricted to the free variables of G and the variables in $fv(range(\sigma))$.

Example 9. Consider the following simple definition for the predicate q

$$q \leftarrow \lambda Z_1. \lambda Z_2. (Z_1 \approx a) \wedge (Z_2 \approx b)$$

that holds only for the tuple (a, b) . Then, consider the goal

$$\leftarrow (R X) \wedge \sim (q X Y)$$

that requests bindings for the variables R, X and Y .

We will underline the selected expression in each step when it is not obvious. We will also omit the substitution in a step if the rule used produces an identity substitution.

0. $\leftarrow \underline{(R X)} \wedge \sim (q X Y)$
1. $\leftarrow \underline{\sim (q X Y)}$
using rule 20.6 and $\theta_1 = \{R/\lambda Z. (Z \approx X) \vee R'\}$
2. $\leftarrow \underline{\sim ((\lambda Z_1. \lambda Z_2. (Z_1 \approx a) \wedge (Z_2 \approx b)) X Y)}$
using rules 20.7, 22.5 and then 21.1
3. $\leftarrow \underline{\sim ((X \approx a) \wedge (Y \approx b))}$
using rules 20.7, 22.5 and then 21.2
4. $\leftarrow \underline{\sim (X \approx a) \vee (X \approx a) \wedge \sim (Y \approx b)}$
using rules 20.7 and then 22.6(d)
5. $\leftarrow \underline{\sim (X \approx a)}$
using rule 20.2

In step 4 the procedure generates two branches. The first one terminates immediately with a primitive inequality. The computed answer is $\sim (X \approx a)$ and $\{R/\lambda Z. (Z \approx X) \vee R'\}$. The second branch will continue as follows:

5. $\leftarrow (X \approx a) \wedge \sim (Y \approx a)$
using rule 20.2
6. $\leftarrow \sim (Y \approx a)$
using rule 20.5 and $\theta_6 = \{X/a\}$

The procedure will now terminate successfully yielding the substitution $\{R/\lambda Z.(Z \approx a) \vee R', X/a\}$ and the primitive goal $\sim (Y \approx a)$.

Implementation

In order to verify and experiment with the ideas presented in this paper, an implementation of the proposed proof procedure has been undertaken. For our implementation we extended an existing interpreter for the language $\mathcal{H}$ introduced in (Charalambidis et al. 2010; 2013). The modified interpreter supports the language $\mathcal{H}_{cn}$ and implements exactly the proof procedure presented in this section. The current version of the interpreter can be retrieved from <https://www.github.com/acharal/hopes>.

The source language⁵ of the interpreter uses an extended Prolog-like syntax, similar to the one we used for motivation purposes in Section 2. Given a source program, the interpreter first derives the types of the predicates using a specialized but relatively straightforward type-inference algorithm. The well-typed program is then transformed to $\mathcal{H}_{cn}$ -like syntax. The interpreter works interactively, i.e., in the same way as ordinary Prolog interpreters. Every time the user specifies a goal, the type information that has been derived from the program is used to typecheck the goal. Subsequently, the goal together with the program is passed to a prover that implements the rules of the proposed proof procedure. The interpreter has been coded in Haskell and it implements the proof procedure as-is, namely it does not apply any low level compilation or any specialized optimizations.

5 The Soundness Theorem

In this section we state the soundness theorem for the proposed proof procedure.

To specify the soundness theorem for the proof procedure, one has to first define the notion of *correct answer*:

Definition 27. Let P be a program and G be a goal. An answer (θ, Q) for $P \cup \{G\}$ is a basic substitution θ and a primitive goal Q .

Definition 28. Let P be a program and let $G = \leftarrow A$ be a goal. An answer $(\theta, \leftarrow A')$ is a *correct answer* for $\text{comp}(P) \cup \{G\}$ if for every model M of $\text{comp}(P)$ and for every state s over the domain of M , $\llbracket A\theta \rrbracket_s(M) \supseteq_o \llbracket A' \rrbracket_s(M)$.

The soundness theorem (Theorem 1 below) essentially amounts to demonstrating that every computed answer is a correct answer. To establish this fact, we need a lemma for each one of Definitions 21, 20 and 22. In particular, Lemma 2 demonstrates the correctness of Definition 21, Lemma 3 shows the correctness of Definition 22, and Lemma 4 the correctness of Definition 20 regarding single-step derivations. Lemma 5 actually shows the correctness of

⁵The language is called *HOPES* standing for Higher-Order Prolog with Extensional Semantics.

many-step derivations. The most demanding proof is that of Lemma 3. The proof of Lemma 3 (as-well-as the proof of the auxiliary Lemma 6 that will be given below) utilize one important and common assumption that exists in all previous works on constructive negation: we assume that “the language contains an infinite number of constants and functions” (see (Chan 1989, page 478)).

Lemma 2. Let P be a program, let E, E' be body expressions of type o and assume that E is reduced to E' . Then, for every model M of $\text{comp}(P)$ and for every state s over the domain of M , it holds $\llbracket E \rrbracket_s(M) = \llbracket E' \rrbracket_s(M)$.

Proof. Straightforward using the semantics of $\mathcal{H}_{cn}$. $\square$

Lemma 3. Let P be a program, let E, E' be body expressions of type o and assume that E is negatively-reduced to E' . Then, for every model M of $\text{comp}(P)$ and for every state s over the domain of M , it holds $\llbracket E \rrbracket_s(M) = \llbracket E' \rrbracket_s(M)$.

Lemma 4. Let P be a program, let $G = \leftarrow A$ and $G' = \leftarrow A'$ be goals and let θ be a basic substitution such that $G \xrightarrow{\theta} G'$. Then, for every model M of $\text{comp}(P)$ and for every state s over the domain of M , it holds $\llbracket A\theta \rrbracket_s(M) \supseteq_o \llbracket A' \rrbracket_s(M)$.

Proof. By case analysis and standard logical arguments. $\square$

Lemma 5. Let P be a program and $G = \leftarrow A$ be a goal. Let $G_0 = G, G_1 = \leftarrow A_1, \dots, G_n = \leftarrow A_n$ be a derivation of length n using basic substitutions $\theta_1, \dots, \theta_n$. Then, for every model M of $\text{comp}(P)$ and for every state s over the domain of M , $\llbracket A\theta_1 \dots \theta_n \rrbracket_s(M) \supseteq \llbracket A_n \rrbracket_s(M)$.

Proof. Using Lemma 4, Lemma 1 and induction on n . $\square$

Finally, we will need another auxiliary lemma, which will be used to obtain the Soundness theorem:

Lemma 6. Let E be a body expression of type o and G a conjunction of primitive inequalities. Then, for all states s and interpretations I if $\llbracket E \rrbracket_s(I) \supseteq_o \llbracket G \rrbracket_s(I)$ then $\llbracket E \rrbracket_s(I) \supseteq_o \llbracket G' \rrbracket_s(I)$ where G' is a conjunction of primitive inequalities that is obtained from G by restricting it to the free variables of E .

Using the above, we obtain the soundness theorem:

Theorem 1 (Soundness Theorem). Let P be a program and G be a goal. Then, every computed answer for $P \cup \{G\}$ is a correct answer for $\text{comp}(P) \cup \{G\}$.

Proof. A direct consequence of Lemma 5 using the auxiliary Lemma 6. $\square$

6 Future Work

We have introduced the higher-order language $\mathcal{H}_{cn}$ which supports constructive negation. There have been various proposals for higher-order logic programming languages, the most notable among them being Hilog (Chen, Kifer, and Warren 1989; 1993) and λ -Prolog (Miller and Nadathur 1986; Nadathur and Miller 1990); to our knowledge however, none of these languages has ever before been extended to support constructive negation.

One possible direction for future work is the investigation of completeness results for the proposed proof-procedure. This is an interesting venue for research but appears to be quite demanding due to the non-trivial nature of the proof procedure (note that, to our knowledge, no completeness results have been reported for the work of (Chan 1989)). Possibly a completeness result could be obtained for the case of finite-tree queries. An easier goal for future research would be to extend the language with other useful programming features that would enhance its higher-order style of programming (such as for example, type polymorphism).

References

- Barendregt, H. P. 1984. *The Lambda Calculus – Its Syntax and Semantics*, volume 103 of *Studies in Logic and the Foundations of Mathematics*. North-Holland.
- Chan, D. 1988. Constructive negation based on the completed database. In *ICLP/SLP*, 111–125.
- Chan, D. 1989. An extension of constructive negation and its application in coroutining. In *NACLP*, 477–493.
- Charalambidis, A.; Handjopoulos, K.; Rondogiannis, P.; and Wadge, W. W. 2010. Extensional higher-order logic programming. In Janhunen, T., and Niemelä, I., eds., *JELIA*, volume 6341 of *Lecture Notes in Computer Science*, 91–103. Springer.
- Charalambidis, A.; Handjopoulos, K.; Rondogiannis, P.; and Wadge, W. W. 2013. Extensional higher-order logic programming. *ACM Transactions on Computational Logic* 14(3).
- Chen, W.; Kifer, M.; and Warren, D. S. 1989. Hilog as a platform for database languages. *IEEE Data Eng. Bull.* 12(3):37–44.
- Chen, W.; Kifer, M.; and Warren, D. S. 1993. Hilog: A foundation for higher-order logic programming. *J. Log. Program.* 15(3):187–230.
- Clark, K. L. 1977. Negation as failure. In *Logic and Data Bases*, 293–322.
- Gelfond, M., and Lifschitz, V. 1988. The stable model semantics for logic programming. In Kowalski, R. A., and Bowen, K. A., eds., *ICLP/SLP*, 1070–1080. MIT Press.
- Lifschitz, V. 2008. What is answer set programming? In Fox, D., and Gomes, C. P., eds., *AAAI*, 1594–1597. AAAI Press.
- Lloyd, J. W. 1987. *Foundations of Logic Programming, 2nd Edition*. Springer.
- Miller, D., and Nadathur, G. 1986. Higher-order logic programming. In Shapiro, E. Y., ed., *ICLP*, volume 225 of *Lecture Notes in Computer Science*, 448–462. Springer.
- Nadathur, G., and Miller, D. 1990. Higher-order horn clauses. *J. ACM* 37(4):777–814.