

Towards a Type Discipline for Answer Set Programming

Camillo Fiorentini, Alberto Momigliano, and Mario Ornaghi

Dipartimento di Scienze dell'Informazione, Università degli Studi di Milano, Italy
{fiorenti,momiglia,ornaghi}@dsi.unimi.it

Abstract. We argue that it is high time that types had a beneficial impact in the field of Answer Set Programming and in particular Disjunctive Datalog as exemplified by the *DLV* system. Things become immediately more challenging, as we wish to present a type system for *DLV-Complex*, an extension of *DLV* with uninterpreted function symbols, external implemented predicates and types. Our type system owes to the seminal polymorphic type system for Prolog introduced by Mycroft and O'Keefe, in the formulation by Lakshman and Reddy. The most innovative part of the paper is developing a *declarative grounding* procedure which is at the same time appropriate for the operational semantics of ASP and able to handle the new features provided by *DLV-Complex*. We discuss the soundness of the procedure and evaluate informally its success in reducing, as expected, the set of ground terms. This yields an automatic reduction in size and numbers of (non isomorphic) models. Similar results could have only been achieved in the current untyped version by careful use of generator predicates in lieu of types.

Keywords: Answer set programming, type checking, grounding, many sorted interpretation.

1 Introduction

The advantages of static type checking for programming languages are almost universally recognized and well-understood, although types managed to make it into logic programming somewhat belatedly [21]. From the pioneering paper [18] advocating the introduction of types in Prolog, different approaches emerged, belonging roughly to two main camps: the descriptive and the prescriptive one. The former aims at capturing some aspects of a program behavior, for example its success set, up to much more complex static and dynamic properties as with CiaoPP [12].

We favor the *prescriptive* approach, where types are an integral part of the program's meaning, thus allowing the user to discard ill-formed elements at compile time. This helps the developer to proceed in a more disciplined way, being able to receive early feedback about mistakes that may be hard to find, especially in a purely declarative setting in the non-infrequent case where the answer would be merely “no”.

The type theory of logic programming has significantly developed in the passing years, following the ever-increasing role of types in the general theory of programming languages – see the type system of Mercury [14] and λ Prolog [19] for a modern take to prescriptive typing. Still, for the sake of this paper, we will adapt to disjunctive

logic programming a rather elementary SML-like polymorphic discipline; in addition we will enrich it semantically, following Lakshman & Reddy’s approach [15], where typed logic programs are interpreted over first order many sorted structures, see also [13].

Our objective here is to advocate the usefulness of prescriptive typing for Answer Set Programming (ASP) [9] in general. ASP is a form of declarative programming based on the stable models semantics [10]. It has been proposed as a language for knowledge representation and as a tool for formalizing and solving hard search problems [1]. The main idea is to reduce a search problem encoded with a (disjunctive) logic program P to the problem of computing stable models of P using an answer set solver, namely a system for generating stable models, such as Smodels [22] and DLV [16]. In this sense ASP’s operational semantics is very different from the usual SLD(NF)-resolution [17]. This is also why typing ASP is interesting and less banal than at first sight: the original paper by Mycroft and O’Keefe was entirely syntactical and so the vast majority of the ensuing research was intrinsically biased towards standard SLD(NF)-resolution.

Although some papers in the ASP literature use a many sorted language or deals with object-oriented features — we refer to Sect. 6 for a detailed discussion — it is safe to say that types have never been *integrated* into ASP; in fact the above papers do not consider a *many-sorted model theory*, nor any form of *parametric* polymorphism. The present paper introduces static type-checking in DLV-Complex [5], an extension of DLV with uninterpreted function symbols – hence we are leaving the comfort of Datalog – as well with external types and predicates requiring an oracle to be computable. In other terms, one may exploit external sources of knowledge. To take into account this possibility, we consider programs over *grounding structures*. The latter provide a declarative semantics for grounding, as they play the role of pre-interpretations [17] in presence of an external implementation of data types and related operations.

The paper is organized as follows: we begin in Section 2 by briefly illustrating the role of typing in ASP through examples. Section 3 and 4 give a static semantics to the language both in terms of types and many-sorted interpretations. In Section 5 we extend the above analysis to DLV-Complex’s programs and offer a rational reconstruction of type-driven grounding. After discussing related work and summing up the results in Section 6, we add a short Appendix (A) explaining the basics of ASP for the reader unfamiliar with this topic.

2 An Informal Introduction to Typed DLV

For the sake of this paper, an ASP program P is a set of rules (or clauses) of the form:

$$A_1 \vee \dots \vee A_n \leftarrow B_1, \dots, B_k, \text{not } B_{k+1}, \dots, \text{not } B_m$$

where $A_1, \dots, A_n, B_1, \dots, B_m$ are atoms, i.e., formulas of the form $p(t)$, where t stands for a sequence t_1, \dots, t_j of terms. Usually *not* is referred to as “negation as failure” or “default negation” and \vee as “epistemic disjunction”.

Example 1. The following pencil problem is a simple non-typed ASP program, which illustrates the meaning of default negation.

```

color(red). color(blue). pencil(p1). pencil(p2).      % facts
nice(X) :- pencil(X), color(Y), not ugly(Y).         % clause c1
ugly(Y) :- color(Y), pencil(X), not nice(X).         % clause c2

```

By rule c1, we can infer that a pencil X is nice if we do not have evidence for $\text{ugly}(Y)$ (i.e., if we can assume $\text{not ugly}(Y)$ as a default). Symmetrically, by rule c2, we can infer $\text{ugly}(Y)$ if we do not have evidence for $\text{nice}(X)$. Let us assume that “not ugly(Y)” holds. We infer that all the pencils are nice, i.e., we have the answer set $\{\text{nice}(p1), \text{nice}(p2)\}$. Symmetrically, assuming “not nice(X)” we obtain the answer set $\{\text{ugly}(red), \text{ugly}(blue)\}$.

As shown by the example, an ASP program P represents a problem that may have multiple solutions, represented by its answer sets. The answer set semantics is defined in terms of *grounding*: let $\text{grnd}(C)$ be the set of the ground instances of a clause C , obtained by substituting its variables with ground terms of the language underlying P . We define $\text{grnd}(P) = \cup_{C \in P} \text{grnd}(C)$. An answer set is a set of ground literals. The answer sets of P are the stable models of $\text{grnd}(P)$ (see the Appendix for the formal definitions).

An ASP solver computes the stable models in two stages: firstly it builds an optimized version of $\text{grnd}(P)$, namely a set $\text{grnd}^\bullet(P)$ of ground clauses with the same semantics of $\text{grnd}(P)$. Then it generates the answer sets of $\text{grnd}^\bullet(P)$.

Example 2. In DLV, the pencil program of Example 1 grounds to:

```

% EDB facts:
color(red). color(blue). pencil(p1). pencil(p2).
% Residual ground instantiation of the program:
ugly(red) :- not nice(p1). ugly(red) :- not nice(p2).
ugly(blue) :- not nice(p1). ugly(blue) :- not nice(p2).
nice(p1) :- not ugly(red). nice(p1) :- not ugly(blue).
nice(p2) :- not ugly(red). nice(p2) :- not ugly(blue).

```

We point out that only 12 simplified clauses of the 36 of $\text{grnd}(\text{pencil})$ are generated, called the *residual ground instantiation*. For example, $\text{ugly}(red) :- \text{color}(red), \text{pencil}(p1), \text{not nice}(p1)$ is simplified to $\text{ugly}(red) :- \text{not nice}(p1)$, because the facts $\text{color}(red)$ and $\text{pencil}(p1)$ are true in every model. In contrast, the clause $\text{ugly}(p1) :- \text{color}(p1), \text{pencil}(p1), \text{not nice}(p1)$ has not been generated because the atom $\text{color}(p1)$ in the body is false in every stable model. We get the answer sets:

```

{color(red), color(blue), pencil(p1), pencil(p2), nice(p1), nice(p2)}
{color(red), color(blue), pencil(p1), pencil(p2), ugly(red), ugly(blue)}

```

Now we come to the main issue of this paper. Looking at the example, one easily recognizes that the optimizations performed by the grounding algorithm, henceforth indicated by *GA*, are related to the implicit presence of two types of objects: pencils and colors. Such types are *interpreted in the same way in all stable models*. We believe that making typing explicit has several advantages. In particular, the programmer has a better control on the grounding process, because a type-driven grounder will not generate clauses corresponding to type information, thus yielding smaller $\text{grnd}^\bullet(P)$ and smaller answer

sets.¹ Furthermore there is a clear-cut distinction between *data base facts*, which represent relevant pieces of information to be shown in answer sets, and *typing constraints*, which only have the role of narrowing grounding, as exemplified next.

Example 3. A typed program is (informally) declared as a unit. We use a concrete syntax inspired by [15]. The typed version of the program of Example 1 is:

```
unit pencils.
type pencil --> p1 ; p2.      % Cf. SML's datatype pencil = p1 | p2;
type color --> red ; blue.
pred nice(pencil), ugly(color).
prog
  nice(X) :- not ugly(Y).
  ugly(X) :- not nice(Y).
```

The *enumerated* type `color` is interpreted as the set of constants $\{\text{red}, \text{blue}\}$, and `pencil` as $\{\text{p1}, \text{p2}\}$. The `pred` keyword introduces a set of predicate symbols and their declared types. By `nice(pencil)` we indicate that `nice` requires an argument of type `pencil` to be a well-typed predicate. Finally, we have the program section, whose clauses are well-typed. Of course we do not need to declare the type of the variables, as they can be inferred. In the example, the rules for clauses in Sect. 5 would yield these well-typing judgments of the form $\Gamma \vdash H \leftarrow B : \text{clause}$.

$$\begin{aligned} X : \text{pencil}, Y : \text{color} &\vdash \text{nice}(X) \leftarrow \text{not ugly}(Y) : \text{clause} \\ X : \text{color}, Y : \text{pencil} &\vdash \text{ugly}(X) \leftarrow \text{not nice}(Y) : \text{clause} \end{aligned}$$

We will explicitly connect the concrete and abstract syntax in Example 6. A typed *GA* will use the variable contexts $X : \text{color}, Y : \text{pencil}$ instead of the predicates `pencil(X)` and `color(Y)` in the untyped program. The difference is that the untyped *GA* catalogues `color(red)`, `color(blue)`, ... as “external data base” (EDB) facts and puts them in the answer sets, as shown in Example 2, since it has no mean to distinguish type and data-base information. In contrast, a typed *GA* would *generate* only the residual grounding instantiation of Example 2, which has the more concise answer sets $\{\text{nice}(\text{p1}), \text{nice}(\text{p2})\}$ and $\{\text{ugly}(\text{red}), \text{ugly}(\text{blue})\}$.

We remark that our model-theory fix the interpretation of external types, functions as well as of equality, but not of program predicates. The latter depends on the program and may have a variety of models. In pure logic programming, types, functions and equality are interpreted on term models. Equality is axiomatized by Clark’s Equality Theory (CET) [17]. For the predicates defined in our example, CET includes the standard equality axioms and the freeness axioms $\neg(\text{red} = \text{blue})$ and $\neg(\text{p1} = \text{p2})$. A set of equations has a solution in CET iff a unifier exists. Thus standard unification is a sound and complete decision method for CET. If we add non-free functions, we no longer have pure term models and the unification algorithm cannot be applied as such.

Example 4. The following program uses the externally implemented type `#int` of integers, where we use the notation `#` to indicate external types and predicates, as in DLV-Complex [5].

¹ We remark that constructions such as `#hide` in *Smodels*, may not show the type information, but the latter is still present during the generation of the model.

```

unit usr.
type #int.
type usr --> john ; ted.
func age: usr -> #int.
    age(john) = 12.
    age(ted) = 15.
pred older(usr,usr,#int).
prog
    older(P1,P2,A) :- age(P1) = age(P2)+A, A > 0.

```

The domain of the external type `#int` is predefined, together with the usual operations on it. In our example, the function `age` is defined by two equations, although it could be an arbitrarily complex computable function. Both `age` and the predefined `+`, `=`, `>` are used in the body of the program rule. An ASP system should be able to solve constraint problems, preferably at the grounding stage. If this were the case, the user would not have to specify a `#maxint` and the `older` clause would not have to be grounded up to that.

The next example uses parametric polymorphism. We adopt the usual Prolog notation for lists `[_|_]` as well as using `'A` for the (quantified) type variable α (following SML's notation).

Example 5. We assume that `#list` is an external type constructor, while `rec` is internal.

```

unit store.
type usr --> john ; ted.
type #int.
type rec('A,'B) --> r('A,'B).
type #list('A).
pred is_user(usr), stored_at(#int,'A), store(#list(rec(#int,'A))).
prog
    is_user(U) :- stored_at(K,U).
    stored_at(K,E) :- store(S), #member(r(K,E),S).
    store([r(1,john),r(2,ted)]).

```

Note that for every instantiation $[T/'A]$, the constructor `#list(T)` yields an external implementation of finite lists with elements of type T , together with a set of pre-defined operations such as `#member('A,#list('A))`.

Since the signature of a DLV-Complex program consists of two distinct parts, one containing functions and (external) data predicates, the other (internal) predicates, we present the type system, the abstract syntax and the intended model semantics in three steps over the next three Sections. We firstly introduce a many-sorted first order data signature Σ^* induced by a polymorphic type system. Then we define the *intended data models* as suitable Σ^* -interpretations. Finally, we introduce the program signature, DLV-Complex programs and their intended models.

3 The Type and Term Language

We start from the type system. Types are defined inductively applying a type constructors $c(\#c)$ to *internal* and *external* base types. Internal atomic types include a unit type

and a constant “o” denoting the type of propositions. To stay inside first-order logic, “o” can only appear at the right of an arrow type. Terms are obtained from logical variables X and function symbols. Signatures give kinds to (any) type constructors and types to internal function symbols and external predicates. The latter two can be used polymorphically.

Base types	T_0	$::= \alpha \mid \mathbf{o} \mid \mathbf{unit} \mid \#int \mid \#string \dots$
Compound types	T_{n+1}	$::= c(\mathbf{T}_n) \mid \#c(\mathbf{T}_n)$
Types	T	$::= \bigcup_{n \geq 0} T_n$
Signature	Σ	$::= \emptyset \mid \Sigma, c : \mathbf{type} \rightarrow \mathbf{type} \mid \Sigma, \#c : \mathbf{type} \rightarrow \mathbf{type}$ $\mid \Sigma, f : \forall \alpha. \mathbf{T} \rightarrow \mathbf{T}' \mid \Sigma, \#p : \forall \alpha. \mathbf{T} \rightarrow \mathbf{o}$
Contexts	Γ	$::= \emptyset \mid \Gamma, \alpha : \mathbf{type} \mid \Gamma, X : T$
Terms	t	$::= X \mid f_{\mathbf{T}}(t)$
Atoms	A	$::= \#p(t) \mid t = t'$

Before we discuss type checking for the above syntactic categories (Fig. 1), some common notation and conventions: tokens in bold such as \mathbf{T} stand for a sequence, i.e. T_1, \dots, T_n . We view Σ and Γ as sets and we use Γ, e for $\Gamma \cup \{e\}$ with $e \notin \Gamma$. We will often write α instead of $\alpha : \mathbf{type}$. $\Gamma \vdash_{\Sigma} J_1, \dots, J_n$ abbreviates $\Gamma \vdash_{\Sigma} J_1, \dots, \Gamma \vdash_{\Sigma} J_n$. We often write $\Gamma \vdash J$, leaving the signature Σ understood. In the rule sf , f is new in Σ . In the rules for Declarations, Types, Terms and Atoms, Γ is assumed to be a well-formed context.

In the rules fd and pd , $[I/\alpha]$ indicates the substitution of the variables α by the types I . It is worth remarking that while in our type systems constants can be parameterized by type *schemata*, rather than closed (simple) types, we do not support *impredicative* polymorphism. Indeed, we can think of this still as form of prenex polymorphism, where type schema can be seen as axiom schema in first-order logic.

Example 6. The concrete syntax of Example 5 corresponds to the signature:

$$\begin{aligned} &usr : \mathbf{type}, \#int : \mathbf{type}, john : usr, ted : usr, 0 : \#int, + : \#int, \#int \rightarrow \#int, \dots \\ &rec : \mathbf{type}, \mathbf{type} \rightarrow \mathbf{type}, r : \forall \alpha, \beta. \alpha, \beta \rightarrow rec(\alpha, \beta), \#list : \mathbf{type} \rightarrow \mathbf{type}, \\ &nil : \forall \alpha. \#list(\alpha), cons : \forall \alpha. \alpha, \#list(\alpha) \rightarrow \#list(\alpha), \#member : \forall \alpha. \alpha, \#list(\alpha) \rightarrow \mathbf{o} \dots \end{aligned}$$

The system in Fig. 1 is an extension with rules for external sources of knowledge of the one presented in [15]. One difference however is in the judgments “ $f_{\mathbf{T}} \text{ decl}$ ” and “ $\#p_{\mathbf{T}} \text{ decl}$ ”, which introduce *indexed* function and predicate declarations. In [15] instead, indexes are only used at the semantic level, to provide a first-order interpretation of polymorphic declarations. Indexes can be reconstructed as we will discuss in Sect. 4.

Some relevant admissible properties of the type system are listed below, where $fv(J)$ indicates the set of the variables occurring free in J and $\Gamma|_{fv(J)}$ the restriction of Γ to terms containing variables in $fv(J)$. Here we have J range over Decl, Typs, Trms and Atms.

- uq) Uniqueness of typing: $\Gamma \vdash_{\Sigma} t : T$ and $\Gamma \vdash_{\Sigma} t : T'$ implies $T = T'$, modulo α -conversion.
- str) Strengthening: if $\Gamma \vdash_{\Sigma} J$ then every variable free in J is declared in Γ . Moreover, $\Gamma|_{fv(J)} \vdash_{\Sigma} J$.

Cntx	$\frac{}{\emptyset \text{ cntx}} \text{cn0}$	$\frac{\Gamma \text{ cntx}}{\Gamma, \alpha \text{ cntx}} \text{cn1}$	$\frac{\Gamma \text{ cntx} \quad \Gamma \vdash_{\Sigma} T : \text{type}}{\Gamma, X : T \text{ cntx}} \text{cn2}$
Sig	$\frac{}{\emptyset \text{ sig}} \text{s0}$	$\frac{\Sigma \text{ sig}}{\Sigma, (c : \text{type} \rightarrow \text{type}) \text{ sig}} \text{sc}$	$\frac{\Sigma \text{ sig}}{\Sigma, (\#c : \text{type} \rightarrow \text{type}) \text{ sig}} \text{sc\#}$
	$\frac{\Sigma \text{ sig} \quad \alpha \vdash_{\Sigma} T, T : \text{type}}{\Sigma, f : \forall \alpha. T \rightarrow T \text{ sig}} \text{sf}$	$\frac{\Sigma \text{ sig} \quad \alpha \vdash_{\Sigma} T : \text{type}}{\Sigma, \#p : \forall \alpha. T \rightarrow \text{o sig}} \text{sp}$	
Decl	$\frac{\Gamma \vdash_{\Sigma} I : \text{type} \quad f : \forall \alpha. T \rightarrow T \in \Sigma}{\Gamma \vdash_{\Sigma} f_I : T[I/\alpha] \rightarrow T[I/\alpha] \text{ decl}} \text{fd}$	$\frac{\Gamma \vdash_{\Sigma} I : \text{type} \quad \#p : \forall \alpha. T \rightarrow \text{o} \in \Sigma}{\Gamma \vdash_{\Sigma} \#p_I : T[I/\alpha] \rightarrow \text{o decl}} \text{pd}$	
Types	$\frac{}{\Gamma \vdash_{\Sigma} \text{unit} : \text{type}} \text{Tu}$	$\frac{}{\Gamma, \alpha \vdash_{\Sigma} \alpha : \text{type}} \text{T0}$	
	$\frac{\Gamma \vdash_{\Sigma} T : \text{type} \quad c : \text{type} \rightarrow \text{type} \in \Sigma}{\Gamma \vdash_{\Sigma} c(T) : \text{type}} \text{Ti}$		
Trms	$\frac{}{\Gamma, X : T \vdash_{\Sigma} X : T} \text{t0}$	$\frac{\Gamma \vdash_{\Sigma} t : T \quad \Gamma \vdash_{\Sigma} f_I : T \rightarrow T \text{ decl}}{\Gamma \vdash_{\Sigma} f_I(t) : T} \text{ti}$	
Atms	$\frac{\Gamma \vdash_{\Sigma} t : T \quad \Gamma \vdash_{\Sigma} \#p_I : T \rightarrow \text{o decl}}{\Gamma \vdash_{\Sigma} \#p(t) : \text{o}} \text{pi}$		$\frac{\Gamma \vdash_{\Sigma} t : T \quad \Gamma \vdash_{\Sigma} t' : T}{\Gamma \vdash_{\Sigma} t = t' : \text{o}} \text{eq}$

Fig. 1. Type checking rules

wk) Weakening: if $\Gamma \subseteq \Gamma'$ and $\Gamma \vdash_{\Sigma} J$, then $\Gamma' \vdash_{\Sigma} J$.

sbi) Substitution: we implicitly assume the α -conversions needed to avoid capture and, for $\theta \equiv [T/\alpha]$, $\Gamma\theta$ is defined as: $\emptyset\theta = \emptyset$, $(\Gamma, \alpha : \text{type})\theta = (\Gamma\theta)$, $\alpha : \text{type}$ and $(\Gamma, X : T')\theta = (\Gamma\theta)$, $X : (T'\theta)$.

$$\frac{\Gamma, \alpha \vdash_{\Sigma} T : \text{type} \quad \Gamma \vdash_{\Sigma} T' : \text{type}}{\Gamma[T'/\alpha] \vdash_{\Sigma} T[T'/\alpha] : \text{type}} \text{sb1} \quad \frac{\Gamma, X : T' \vdash_{\Sigma} t : T \quad \Gamma \vdash_{\Sigma} t' : T'}{\Gamma \vdash_{\Sigma} t[t'/X] : T} \text{sb2}$$

4 Grounding Structures

We firstly introduce a many-sorted first order signature Σ^* . Then we define grounding structures as suitable Σ^* -interpretations. A polymorphic signature Σ yields an infinite many sorted signature Σ^* in the following way:

- Every ground type T is a sort in Σ^* ;
- for every $f : \forall \alpha. T \rightarrow T$ such that $\emptyset \vdash_{\Sigma} f_I : T[I/\alpha] \rightarrow T[I/\alpha] \text{ decl}$ is provable, $f_I : T[I/\alpha] \rightarrow T[I/\alpha]$ is a function declaration in Σ^* ;

- for every $\#p : \forall \alpha. \mathbf{T} \rightarrow \circ$ such that $\emptyset \vdash_{\Sigma} \#p_{\mathbf{I}} : \mathbf{T}[\mathbf{I}/\alpha] \rightarrow \circ$ decl is provable, $\#p_{\mathbf{I}} : \mathbf{T}[\mathbf{I}/\alpha] \rightarrow \circ$ is a predicate declaration in Σ^* .

Σ^* -interpretations are defined in a way similar to [13]:

Definition 1. A Σ^* -interpretation is a function ι such that:

- Every T in Σ^* is mapped to a set $\iota(T)$, called the domain of T .
- Every $f_{\mathbf{I}} : \mathbf{T} \rightarrow T$ in Σ^* is mapped to a function $\iota(f_{\mathbf{I}}) : \iota(\mathbf{T}) \rightarrow \iota(T)$.
- Every $\#p_{\mathbf{I}} : \mathbf{T} \rightarrow \circ$ in Σ^* is mapped to a relation $\iota(\#p_{\mathbf{I}}) \subseteq \iota(\mathbf{T})$.

We are interested in a fixed interpretation G , called *grounding structure* or *pre-interpretation*.² The grounding structure of a program is defined by its type and function sections and the predefined types thereby declared.

Example 7. Recall the signature Σ of Example 5. The intended grounding structure is the interpretation G defined as follows.

- Base type usr : $G(usr) = \{john, ted\}$, $G(john) = john$, $G(ted) = ted$.
- Base type $\#int$: it is interpreted according to its external implementation.
- Level 1 type $rec(\#int, usr)$: $G(rec(\#int, usr))$ is the set generated by $r_{\#int, usr}$ starting from the interpretations of level 0 $G(\#int)$ and $G(usr)$. A representation is:
 - $G(rec(\#int, usr)) = \{r_{\#int, usr}\} \times G(\#int) \times G(usr)$;
 - $G(r_{\#int, usr})$ is the map $\langle i \in G(\#int), u \in G(usr) \rangle \mapsto \langle r_{\#int, usr}, i, u \rangle$.
 By abuse of notation, we represent triples $\langle r_{\#int, usr}, i, u \rangle$ as “terms” $r(i, u)$.
- Level 1 type $\#list(usr)$: $G(\#list(usr))$ is the structure of the finite lists with elements from $G(usr)$. The predefined operations are interpreted according to their implementation.
- ... and so on. We remark that the set-theoretic representation applies at each level.

We now move to the evaluation of a term with respect to an *assignment* in an interpretation.

Definition 2. Let ι be a Σ^* -interpretation and Γ a context. An assignment a for Γ in ι maps every β of Γ to a sort $a(\beta)$ of Σ^* and every $X : T$ of Γ to a value $v \in \iota(Ta)$, where Ta is the sort of Σ^* obtained by grounding each variable β occurring in T by the sort $a(\beta)$.

Given ι , we denote by $ass(\Gamma, \iota)$ the set of all the assignments for Γ in ι . Take $\Gamma \vdash_{\Sigma} t : T$ and $a \in ass(\Gamma, \iota)$. The evaluation of t with respect to a , denoted $t \rightsquigarrow_a v$, is performed according to the following rules:

$$\frac{}{X \rightsquigarrow_a a(X)} (ev1) \qquad \frac{t \rightsquigarrow_a v \quad \Gamma_{\Sigma} \vdash f_{\mathbf{I}} : \mathbf{T} \rightarrow T \text{ decl}}{f_{\mathbf{I}}(t) \rightsquigarrow_a \iota(f_{\mathbf{I}a})(v)} (ev2)$$

Theorem 1. Let Σ be a signature, ι a Σ^* -interpretation, $a \in ass(\Gamma, \iota)$ and $\Gamma \vdash_{\Sigma} t : T$. Then there is a unique $v \in \iota(Ta)$ such that $t \rightsquigarrow_a v$.

² Our notion of pre-interpretation generalizes Lloyd’s [17] interpreting #-predicates.

Theorem 1 follows by the existence and uniqueness of the $\Gamma \vdash_{\Sigma} f_I : T \rightarrow T$ decl judgment and as such depends on the indexing of function symbols via a simple table look up “modulo matching” over Σ .

Example 8. Let Σ be the signature of Ex. 5 and $\Gamma = \{\alpha, \beta : \text{type}, X : \alpha, Y : \beta\}$. We have that $\Gamma \vdash_{\Sigma} r(X, Y) : \text{rec}(\alpha, \beta)$. Let G be the interpretation defined in Ex. 7 and $a \in \text{ass}(\Gamma, G)$ such that $a(\alpha) = \#int$, $a(\beta) = usr$, $a(X) = 5$ and $a(Y) = ted$. Then $r(X, Y) \rightsquigarrow_a \langle r_{\#int,usr}, 5, ted \rangle$.

Finally, the truth relation $\iota \models_a A$ for a well-typed atom A in Γ and $a \in \text{ass}(\Gamma, \iota)$ is defined by cases on A :

- $\iota \models_a \#p_I(t)$ iff $t \rightsquigarrow_a v$ and $v \in \iota(p_I)$;
- $\iota \models_a t_1 = t_2$ iff $t_j \rightsquigarrow_a v$ for $j = 1, 2$.

5 Program Grounding

A *program signature* is of the form $\Sigma \cup \Pi$ where Π is a set of *program predicate* declarations. We remark that Σ still accounts for the external predicates $\#p$. As we saw in Section 2, the latter are interpreted on top of the grounding structure G defined in the type and function section of the program at hand. To formalize this, we resort to G -expansions. We start with $\Sigma \cup \Pi$ as above: the associated many sorted signature $(\Sigma \cup \Pi)^*$ has the form $\Sigma^* \cup \Pi^*$, where Π^* contains the ground predicate declarations $p_I : T \rightarrow o$. A grounding structure G is a Σ^* -interpretation and a G -expansion is a $(\Sigma^* \cup \Pi^*)$ -interpretation ι such that $\iota|_{\Sigma^*} = G$, i.e. the *reduct* of ι w.r.t. Σ^* . We can represent ι as the set

$$H_{\iota} = \{ \langle p_I, v \rangle \mid p_I : T \rightarrow o \in \Pi^* \text{ and } v \in \iota(p_I) \}$$

By abuse of notation, we write $p(v)$ instead of $\langle p_I, v \rangle$ and call it a G -atom. We dub the set of G -atoms the G -base, playing a role analogous to the Herbrand base. In this way, H_{ι} can be thought as a kind of “ G -answer set”.

Example 9. Back to program of Example 5 and the related grounding structure of Example 7. The following interpretation ι is the minimum model of the program (the unit “store”), seen as a set of G -atoms:

$$\{ \text{store}([r(1, john), r(2, ted)]), \text{stored_at}(1, john), \text{stored_at}(2, ted), \text{is_user}(john), \text{is_user}(ted) \}$$

The loose notation $\text{store}([r(1, john), r(2, ted)]), \text{stored_at}(1, john), \dots$ means that

$$\langle \text{cons}, \langle r, 1, john, \rangle, \langle \text{cons}, \langle r, 2, ted \rangle, nil \rangle \rangle \in G(\text{store}_{usr}), \langle 1, john \rangle \in G(\text{stored_at}_{usr}), \dots$$

where $\text{cons} = \text{cons}_{\text{rec}(\#int,usr)}$ and $nil = nil_{\text{rec}(\#int,usr)}$.

Now we can formally introduce programs and enrich the syntax judgments of the previous section. The form of a clause is $A_1 \vee \dots \vee A_n \leftarrow B_1, \dots, B_m$, where A_1, \dots, A_n are program atoms and B_1, \dots, B_m are literals, i.e., atoms, external predicates, negated atoms or negated external predicates. For the sake of simplicity, we do not consider here classical negation, see the appendix.

- Atom $A ::= \dots$ (as before).
- Program Atom $PA ::= p(\mathbf{t})$, with p defined in Π .
- Negated Atom $NA ::= \text{not } A \mid \text{not } PA$.
- Literal $B ::= A \mid PA \mid NA$.
- Head $HD ::= A_1 \vee \dots \vee A_n$; for $n = 0$, $HD \equiv \text{false}$.
- Body $BD ::= B_1, \dots, B_m$; for $m = 0$, $BD \equiv \text{true}$.
- Clause $C ::= HD \leftarrow BD$.

To distinguish external and internal program predicates, we split the judgment for well-formed atoms $p(\mathbf{t}) : \mathfrak{o}^+$ and $\#q(\mathbf{t}) : \mathfrak{o}^\#$, where q is declared in Σ and p in Π . The corresponding rules are omitted, since they correspondingly split the rule pi in Fig. 1. We introduce the judgments $NA : \mathfrak{o}^-$, $HD : \text{head}$, $BD : \text{body}$ and $C : \text{clause}$, where in $B_i : \mathfrak{o}^{s_i}$, the superscript s_i may be $+$, $-$, $\#$.

$$\begin{array}{c}
\frac{\Gamma \vdash_{\Sigma \cup \Pi} A : \mathfrak{o}^+}{\Gamma \vdash_{\Sigma \cup \Pi} \text{not } A : \mathfrak{o}^-} \text{at}^- \qquad \frac{\Gamma \vdash_{\Sigma \cup \Pi} A : \mathfrak{o}^\#}{\Gamma \vdash_{\Sigma \cup \Pi} \text{not } A : \mathfrak{o}^-} \text{at}^- \\
\frac{\Gamma \vdash_{\Sigma \cup \Pi} A_1 : \mathfrak{o}^+, \dots, A_n : \mathfrak{o}^+}{\Gamma \vdash_{\Sigma \cup \Pi} A_1 \vee \dots \vee A_n : \text{head}} h \qquad \frac{\Gamma \vdash_{\Sigma \cup \Pi} B_1 : \mathfrak{o}^{s_1}, \dots, B_m : \mathfrak{o}^{s_m}}{\Gamma \vdash_{\Sigma \cup \Pi} B_1, \dots, B_m : \text{body}} b \\
\frac{\Gamma \vdash_{\Sigma \cup \Pi} HD : \text{head} \quad \Gamma \vdash_{\Sigma \cup \Pi} BD : \text{body}}{\Gamma \vdash_{\Sigma \cup \Pi} HD \leftarrow BD : \text{clause}} cl
\end{array}$$

Let $\Sigma \cup \Pi$ be a program signature, G a grounding structure and \mathfrak{t} be a G -expansion. The truth relation $\mathfrak{t} \models_a A$ is extended to clauses as usual. Similarly, the standard properties of the type system stated on Page 122 also hold.

Now, we aim to define the grounding of a program P with respect to the structure G , notation $\text{grnd}_G(P)$, where P consists of a set of clauses $\forall \alpha : \text{type}. \forall \mathbf{X} : \mathbf{T}. C$ such that $\alpha : \text{type}, \mathbf{X} : \mathbf{T} \vdash C : \text{clause}$. Let $a \in \text{ass}(\Gamma, G)$; if we write $\Gamma \vdash C$, for “ $\Gamma \vdash_G C : \text{clause}$ ”, we can define grounding on the structure of C :

$$\begin{aligned}
\text{grnd}_G(\Gamma \vdash p(\mathbf{t}), a) &= p(\mathbf{v}) \quad \text{where } \mathbf{t} \rightsquigarrow_a \mathbf{v} \\
\text{grnd}_G(\Gamma \vdash \text{not } A, a) &= \text{not } \text{grnd}_G(\Gamma \vdash A, a) \\
\text{grnd}_G(\Gamma \vdash B_1, \dots, B_m, a) &= \text{grnd}_G(\Gamma \vdash B_1, a), \dots, \text{grnd}_G(\Gamma \vdash B_m, a) \\
\text{grnd}_G(\Gamma \vdash A_1 \vee \dots \vee A_n, a) &= \text{grnd}_G(\Gamma \vdash A_1, a) \vee \dots \vee \text{grnd}_G(\Gamma \vdash A_n, a) \\
\text{grnd}_G(\Gamma \vdash H \leftarrow B, a) &= \text{grnd}_G(\Gamma \vdash H, a) \leftarrow \text{grnd}_G(\Gamma \vdash B, a) \\
\text{grnd}_G(\Gamma \vdash C) &= \bigcup_{a \in \text{ass}(\Gamma, G)} \{ \text{grnd}_G(\Gamma \vdash C, a) \} \\
\text{grnd}_G(P) &= \bigcup_{\Gamma \vdash C \in P} \text{grnd}_G(\Gamma \vdash C)
\end{aligned}$$

We can see $\text{grnd}_G(P)$ as a set of propositional clauses over the G -base. We have

$$H_{\mathfrak{t}} \models_{\text{prop}} \text{grnd}_G(P) \quad \text{iff} \quad \mathfrak{t} \models P \tag{1}$$

where \models_{prop} is propositional truth. By the equivalence (1), we can define the G -answer sets of P as those of $\text{grnd}_G(P)$. The latter are defined using the Gelfond-Lifschitz transformation, as explained in the appendix.

We now discuss the interaction between grounding and typing: this brings about some immediate problems. The first one is *ground type reconstruction*, as illustrated by the following example.

Example 10

```

pred p(#list('A)), q(#list('A)), r(#list(#char)), t(#nat).
prog
  p([X|Y]) :- q(Y).
  q(Z) :- r(Z).
  r([]).
  t(3).

```

Now, given the ground instance $p([3]) \leftarrow q([])$, the atom $p([3])$ is in the minimal model of the program, which clearly violates the typing discipline. To overcome this problem, we have to assign to $[]$ the “right ground type” T . A solution is to assign an index to constants, apparently similar to how constants are declared in the simply-typed λ -calculus to recover a principal type, viz. $inl_B : A \rightarrow A \vee B$. In our case, $\emptyset \vdash_{\Sigma} []_{\#int} : T$ has solution $T = \#list(\#int)$. Using indexed terms, the above clauses become $p([3]) \leftarrow q([]_{\#int})$, as well as $q([]_{\#char}) \leftarrow r([]_{\#char})$ and $r([]_{\#char})$. Now the standard ASP semantics correctly works.

The reader may wonder whether function indexing can be avoided and thus reconstructed. Interestingly enough, the dual of the property of *transparency* [11], for which the type declaration of every (internal) function is such that every type variable occurring in the domain also occurs in the range, ensures, in the absence of overloading, index reconstruction. Space prevents us from providing a formal proof.

Any non-trivial type signature is a source of infinity. We draw our inspiration from the way in which [5] deals with computable functions. We view type instantiation and grounding of individual variables as two separate phases. The intuition is that polymorphic predicate definitions behave as *open units*, i.e., groups of clauses *abstracting* type variables. During grounding an open unit is instantiated by grounding the type variables over a finite number of ground types of its signature ($tinst(P)$). Subsequently $grnd^{\bullet}(tinst(P))$ grounds individual variables. In the absence of a satisfactory notion of modules and their instantiation in the ASP framework yet, we will not formally define $tinst(P)$. Rather, we directly illustrate a case where modularity and type instantiation do not work properly. In this situation, the user has to provide the type instantiation with a so-called “@with-directive”.

Example 11

```

unit choices.
#maxint=1.
type pencil --> p1 ; p2.
type color --> red ; blue.
func col : pencil -> color.
  col(p1) = red.
  col(p2) = blue.
pred choose(pencil), ok(#int), nice('A), ugly('B).
prog

```

```

nice(X) ∨ ugly(Y).
choose(X) :- nice(X), not ugly(col(X)).
@with ugly(#int):
ok(X+1) :- nice(X), X < 1.

```

Type inference yields:

```

poly:   α, β : type, X : α, Y : β ⊢ niceα(X) ∨ uglyβ(Y) : head
used by: i) X : pencil ⊢ choose(X) ← nicepencil(X), not uglycolor(col(X)) : clause
        ii) X : #int ⊢ ok(X + 1) ← nice#int(X), X < 1 : clause

```

The type instantiation inferred from i) is $[pencil/\alpha, color/\beta]$, while ii) gives the partial instantiation $[\#int/\alpha]$. The `@with`-directive provides $[\#int/\beta]$. The result of the type instantiation is:

```

c1   X : pencil, Y : color ⊢ nice(X) ∨ ugly(Y) : head
c2           X, Y : #int ⊢ nice(X) ∨ ugly(Y) : head
c3           X : pencil ⊢ choose(X) ← nice(X), not ugly(col(X)) : clause
c4           X : #int ⊢ ok(X + 1) ← nice(X), X < 1 : clause

```

The inferred contexts will be relevant next, in Example 13. With the directive `@with ugly(#int)`, the instantiations of *nice* and of *ugly* in c_1 and c_2 do not interfere, i.e. composition is modular. However, the directive `@with ugly(color)` would have set $[color/\beta]$. Then the clauses c_1 and c_2 would have shared the same $ugly_{color}$ predicate.

A sufficient condition to avoid interference is to impose that the head predicates have different indexed names in different type instantiations. In our example, we could obtain this by declaring a unit `allNiceOrAllUgly(α, β)` with $nice, ugly : \forall \alpha, \beta. \alpha \rightarrow o$.

Now we explain the grounding phase. We extend the notion of context so that it contains everything related to the external predicates and operations, while the right-hand side of \vdash contains only the program and the free functions. Namely we transform a clause $\Gamma \vdash C$ into an *extended clause* $\Gamma' | \Delta \vdash C'$, where $\Gamma \subseteq \Gamma'$, Δ is a set of Σ -formulas, which we call *constraints*, and C' now contains only predicates from Π and function symbols that are internal. This can be done by moving external predicates in the constraint context Δ and by introducing suitable equalities in the usual CLP's way.

Example 12. The clauses c_1, \dots, c_4 of Example 11 are transformed into the following extended clauses e_1, \dots, e_4 :

```

e1           X : pencil, Y : color | ⊤ ⊢ nice(X) ∨ ugly(Y)
e2           X, Y : #int | ⊤ ⊢ nice(X) ∨ ugly(Y)
e3   X : pencil, U : color | U = col(X) ⊢ choose(X) ← nice(X), not ugly(U)
e4           X, J : #int | J = X + 1, X < 1 ⊢ ok(J) ← nice(X)

```

In e_3 , the term $col(X)$ occurring in c_3 is replaced by a new special variable $U : color$ and $U = col(X)$ is inserted in the constraints context. Similarly for $J = X + 1$ in e_4 . Furthermore, the external predicate $X < 1$ has been moved to the constraints context.

Finally, we explain how GA uses extended contexts. Take the signature $\Sigma \cup \Pi$ and extended clause $\Gamma | \Delta \vdash C$. Clearly, GA may access both the signature and the grounding structure G . Take an assignment a such that $a \in \text{ass}(\Gamma, G)$. Since Δ contains Σ -formulas and G is a Σ -interpretation, GA can evaluate the truth relation $G \models_a \Delta$.

$$\begin{aligned} \text{egrnd}_G(\Gamma | \Delta \vdash C, a) &= \begin{cases} \text{grnd}_G(\Gamma \vdash C, a) & \text{if } G \models_a \Delta \\ \top & \text{otherwise} \end{cases} \\ \text{egrnd}_G(\Gamma | \Delta \vdash C) &= \bigcup_{a \in \text{ass}(\Gamma, G)} \{ \text{egrnd}_G(\Gamma | \Delta \vdash C, a) \} \end{aligned}$$

It has the following property. Let P be a $(\Sigma \cup \Pi)$ -program and P_e be the corresponding extended program. The grounding $\text{egrnd}_G(P_e)$ is the union of $\text{egrnd}_G(\Gamma | \Delta \vdash C)$, for $\Gamma | \Delta \vdash C \in P_e$. One can prove that the stable models of $\text{grnd}_G(P)$ coincide with those of $\text{egrnd}_G(P_e)$.

Example 13. Take the clause e_3 of Example 12. An assignment for its context is $a_1 = [p1/X, red/U]$. Since $G \models_{a_1} U = col(X)$, we have

$$\text{egrnd}_G(e_3, a_1) = \text{choose}(p1) \leftarrow \text{nice}(p1), \text{not ugly}(red)$$

In contrast, for the assignment $a_2 = [p1/X, blue/U]$, we have $G \not\models_{a_2} U = col(X)$. Thus $\text{egrnd}_G(e_3, a_2) = \top$, which can be ignored. Since we have finite types, we generate all the possible assignments and, at the end of the process, we obtain:

$$\begin{aligned} \text{nice}(p1) \vee \text{ugly}(red) & \quad \text{nice}(p1) \vee \text{ugly}(blue) \\ \text{nice}(p2) \vee \text{ugly}(red) & \quad \text{nice}(p2) \vee \text{ugly}(blue) \\ \text{nice}(0) \vee \text{ugly}(0) & \quad \text{nice}(0) \vee \text{ugly}(1) \\ \text{nice}(1) \vee \text{ugly}(0) & \quad \text{nice}(1) \vee \text{ugly}(1) \\ \text{choose}(p1) \leftarrow \text{nice}(p1), \text{not ugly}(red) & \\ \text{choose}(p2) \leftarrow \text{nice}(p2), \text{not ugly}(blue) & \\ \text{ok}(1) \leftarrow \text{nice}(0) & \end{aligned}$$

As the example shows, a type-driven grounder GA has the following features:

- In the grounding process, only well-typed instances are generated, viz. in Example 13, $\text{nice}(p1) \vee \text{ugly}(p1)$, $\text{nice}(p1) \vee \text{ugly}(p2)$, ... are excluded.
- If a variable X functionally depends on other variables Y , one value is generated for it, depending on the values generated for Y . For instance, in clause e_3 of Ex. 12 the variable U functionally depends on X , thus in Ex. 13 only the clauses $\text{choose}(p1) \leftarrow \text{nice}(p1), \text{not ugly}(red)$ and $\text{choose}(p2) \leftarrow \text{nice}(p2), \text{not ugly}(blue)$ are added.

This should facilitate the generation of small ground instances. More importantly, types give the user a better control of the grounding process.

In the previous example, the grounding structure is finite. To complete the picture, we should consider the general case of possibly infinite grounding structures. We will not introduce new ideas, but consider how types could be managed by existing approaches. Specifically, we consider intelligent grounding w.r.t. the class of finitely ground programs [6], denoted by \mathcal{FG} . Here we explain the idea informally. Intelligent grounding

is performed starting from a modular decomposition $\langle M_1, M_2, \dots, M_n \rangle$ of the program P . The module M_1 does not depend on any other modules and grounds to a G_1 having the same stable models as M_1 . We use G_1 to build a $G_2 = ig(M_2, G_1)$, having the same stable models of $M_1 \cup M_2$. We proceed in this way until we obtain $G_n = ig(M_n, G_{n-1})$, which has the same models of P . If G_n is finite, one says that P is an \mathcal{FG} -program.

Example 14. This contrived example shows a program that can be divided in two modules (we use a, b, \dots as constants of type $\#char$).

```
unit show.
type #int, #char, #list('A).
pred select(#list('A)), show(#list('A)).
prog
  select([a]) v select([4]).                %module m1
  select([b,a]).
  select([X]) :- select([X+X]).
  show(M) :- select(L),select(M), #reverse(L,M). %module m2
```

The module decomposition is $\langle m1, m2 \rangle$. Note that $m2$ is polymorphic. Type inferences yield the following extended clauses:

$$\begin{array}{ll}
l_1 & \vdash \text{select}([a]) \vee \text{select}([4]) \\
l_2 & \vdash \text{select}([b,a]) \\
l_3 & X, I : \#int \mid I = X + X \vdash \text{select}([X]) \leftarrow \text{select}([I]) \\
l_4 & L, M : \#list(\#char) \mid \#reverse(L, M) \vdash \text{show}(M) \leftarrow \text{select}(L), \text{select}(M) \\
l_5 & L, M : \#list(\#int) \mid \#reverse(L, M) \vdash \text{show}(M) \leftarrow \text{select}(L), \text{select}(M)
\end{array}$$

Now we proceed to the intelligent grounding phase, by applying the instantiation algorithm given in [6] and using extended clauses. We firstly ground the module $m1$. We assume that the grounding algorithm uses a, b, c, \dots as the internal representation for characters, $0, 1, 2, \dots$ for numbers and $[v_1, \dots, v_n]$ for lists, so that we do not need to distinguish between values and ground terms. Following [6], we start from the already ground clauses:

$$\vdash \text{select}([a]) \vee \text{select}([4]) \quad \vdash \text{select}([b,a])$$

We consider the set \mathcal{G} of atoms $\text{select}([a])$, $\text{select}([4])$, $\text{select}([b,a])$ occurring in the head of the ground clauses of $m1$. We use \mathcal{G} to instantiate l_3 by looking for all the grounding substitutions that unify $\text{select}([I])$ with one of the atoms in it. We get:

$$X : \#int \mid 4 = X + X \vdash \text{select}([X]) \leftarrow \text{select}([4])$$

Differently from Example 13, the GA tries to instantiate clauses looking for body atoms that match with atoms in the current \mathcal{G} . In this way, the variable X is not instantiated. It has to be grounded by solving the constraint $4 = X + X$. The solution is of course $X = 2$. The result of this step is:

$$\begin{array}{ll}
& \vdash \text{select}([a]) \vee \text{select}([4]) \\
& \vdash \text{select}([b,a]) \\
& \vdash \text{select}([2]) \leftarrow \text{select}([4]) \\
X, I : \#int \mid I = X + X & \vdash \text{select}([X]) \leftarrow \text{select}([I])
\end{array}$$

and the new \mathcal{G} contains also $select([2])$. We iteratively apply the same process, which yields the instantiation $I = 2$ and the equation $2 = X + X$. We get:

$$\begin{array}{l}
\vdash select([a]) \vee select([4]) \\
\vdash select([b,a]) \\
\vdash select([2]) \leftarrow select([4]) \\
\vdash select([1]) \leftarrow select([2]) \\
X, I : \#int \mid I = X + X \vdash select([X]) \leftarrow select([I])
\end{array}$$

Now no new atom is added to \mathcal{G} and we stop the grounding process for $m1$. The next step is to use $m1$ to ground $m2$. We just consider one of the grounding operations activated in this phase, that is the instantiation and simplification of $select(L)$, $select(M)$ by $select([b,a])$ in the first clause of $m2$. We get $\dots \mid \#reverse([b,a], [b,a]) \vdash show([b,a])$, where the body has been simplified because $select([b,a])$ is true. The evaluation of $\#reverse([b,a], [b,a])$ fails and the result is \top , i.e., no new clause is generated.

The main differences with respect to [6] are:

- Typing does not even consider ill typed substitutions.
- The treatment of external functions is shifted to the context. The context also contains the external atoms. Thus, the problem of grounding a clause in presence of predicates and operations defined in the grounding structure G is reduced to a constraint solving problem over G . The class of the groundable programs depends both on the grounding strategy and on the class of the constraints that the grounder can solve. For example, we could reconstruct value invention as treated in [5] in terms of programs that guarantee the solvability of a special class of constraints.

6 Related Work and Conclusions

The introduction of polymorphism in logic programming has required some care: there are predicate definitions that are semantically non-problematic, yet may lead to runtime errors [15, 11]. A sufficient condition is *definitional genericity*, which requires that the type of a defining occurrence of a predicate must be a renaming of the assigned type signature of the predicate. This ensures the type soundness result of [18], stating that if a program and a goal are well-typed, then at each resolution step variables can only be instantiated to terms consistent with their typing.

Although our system owes its static semantics to Mycroft and O’Keefe’s seminal paper, the very different operational semantics of ASP allows us to dispense of those restrictions – as a matter of fact a dual property to transparency is instrumental in inferring type indexes that ensure ground type reconstruction.

Some papers in the ASP literature, e.g. [9, 2] use a many sorted language or deals with object-oriented features. Sorts, however, are merely syntactic sugar, to be translated back into first-order logic, hence they are not given a semantics in terms of many sorted first order signatures, not (polymorphic) type checking is addressed. W.r.t. object orientation OntoDLV [8] has recently extended DLV-Complex, catering to the specification and reasoning on enterprise ontologies; it extends ASP with object-oriented constructs, such as classes, objects, multiple inheritance, sets and lists. The semantics

is based on a notion of well-defined interpretation of a program, which is defined over the Herbrand Universe of the program. Differently from our approach, it lacks a formal notion of “external” structure where to interpret external predicates, moreover “generics” are not considered.

Modern prescriptive types system for logic programming have significantly evolved since the Eighties: for example, Mercury’s type system [14] is a fairly complex extension of [18], supporting higher order types as λ Prolog, but also type classes similar to Haskell, in addition with a form of existential types. Moreover its mode system [20] provides support for subtypes as well as for uniqueness, similar to linear types. We plan to investigate how modes can be used in our approach to deal with grounding, as we have touched upon in Example 13.

Finally we mention the extensive research about polymorphic type inference in relation algebra and (object-oriented) databases, see [7, 4] for examples. However, this is only loosely related to our efforts, as there the aim is to extend a ML-like type system to capture the *principal* type of a program involving database operations or more in general of an expression in relational algebra, where a type is seen as a set of attribute names for a given schema.

In conclusion, we have presented a polymorphic types system for DLV-Complex and developed a declarative semantics for grounding, while studied its impact on model generation. Although our research is at an early stage, we believe that it shows the potential of types in ASP. A fortiori this should hold true for ASP systems such as Smodels [22] that implements a simpler logic. Beside the well-known advantages of prescriptive typing, in the ASP setting one may reap significant benefits w.r.t. the size of grounding; finally, we have shown that certain constraints can be directly codified in terms of typing, further reducing the size of the search space. This fits with current research aimed at finding a tighter relationship between ASP and constraint solving [2]. Directions for future research include:

- quantitatively evaluate the efficiency of type-driven grounding in terms of size and number of non-isomorphic models;
- extend the type system with other constructs, such as aggregate types, subtyping, etc., in order to account of other features of ASP, e.g. *preference* and *cardinality* constraints [16]. In particular, we speculate that we can elegantly define aggregate types using concrete data as in [3];
- From a practical standpoint, evaluate whether it is feasible to decompile the many-sorted language into untyped first order logic so that the present systems can readily be used or it is more effective to redesign grounding to take into full account the typing discipline.

Acknowledgments. This research was partially supported by the MIUR project “Potenziamento e Applicazioni della Programmazione Logica Disgiuntiva”. We wish to thanks the anonymous referees for having appreciated the potential of the paper, notwithstanding several lacunae and imperfections. Those, we have done our best to fix.

References

1. Baral, C.: Knowledge Representation, Reasoning and Declarative Problem Solving. CUP (2003)
2. Baselice, S., Bonatti, P.A., Gelfond, M.: Towards an integration of answer set and constraint solving. In: Gabbriellini, M., Gupta, G. (eds.) ICLP 2005. LNCS, vol. 3668, pp. 52–66. Springer, Heidelberg (2005)
3. Bertoni, A., Mauri, G., Miglioli, P.: A characterization of abstract data as model-theoretic invariants. In: Maurer, H.A. (ed.) ICALP 1979. LNCS, vol. 71, pp. 26–37. Springer, Heidelberg (1979)
4. Buneman, P., Ogori, A.: Polymorphism and type inference in database programming. *ACM Trans. Database Syst.* 21(1), 30–76 (1996)
5. Calimeri, F., Cozza, S., Ianni, G.: External sources of knowledge and value invention in logic programming. *Ann. Math. Artif. Intell.* 50(3-4), 333–361 (2007)
6. Calimeri, F., Cozza, S., Ianni, G., Leone, N.: Computable functions in ASP: Theory and implementation. In: de la Banda, M.G., Pontelli, E. (eds.) ICLP 2008. LNCS, vol. 5366, pp. 407–424. Springer, Heidelberg (2008)
7. Van den Bussche, J., Waller, E.: Polymorphic type inference for the relational algebra. *J. Comput. Syst. Sci.* 64(3), 694–718 (2002)
8. Ricca, F., Gallucci, L., Schindlauer, R., Dell’Armi, T., Grasso, G., Leone, N.: OntoDLV: an ASP-based system for enterprise ontologies. *Journal of Logic and Computation* (to appear, 2009)
9. Gelfond, M.: Answer sets. In: van Harmelen, F., Lifschitz, V., Porter, B. (eds.) *Handbook of knowledge representation*, vol. 7. Elsevier, Amsterdam (2007)
10. Gelfond, M., Lifschitz, V.: The stable model semantics for logic programming. In: ICLP/SLP, pp. 1070–1080 (1988)
11. Hanus, M.: Horn clause programs with polymorphic types: Semantics and resolution. *Theor. Comput. Sci.* 89(1), 63–106 (1991)
12. Hermenegildo, M.V., Puebla, G., Bueno, F., López-García, P.: Integrated program debugging, verification, and optimization using abstract interpretation (and the Ciao system preprocessor). *Sci. Comput. Programming* 58(1-2), 115–140 (2005)
13. Hill, P.M., Topor, R.W.: A semantics for typed logic programs. In: *Types in Logic Programming*, pp. 1–62. MIT Press, Cambridge (1992)
14. Jefferey, D.: Expressive Type Systems for Logic Programming Languages. PhD thesis, The University of Melbourne (2002)
15. Lakshman, T.L., Reddy, U.S.: Typed PROLOG: A semantic reconstruction of the Mycroft-O’Keefe type system. In: ISLP, pp. 202–217 (1991)
16. Leone, N., et al.: The DLV system for knowledge representation and reasoning. *ACM TOCL* 7(3), 499–562 (2006)
17. Lloyd, J.W.: *Foundations of logic programming* (2nd extended ed.). Springer, Heidelberg (1987)
18. Mycroft, A., O’Keefe, R.A.: A polymorphic type system for PROLOG. *Artif. Intell.* 23(3), 295–307 (1984)
19. Nadathur, G., Qi, X.: Optimizing the runtime processing of types in polymorphic logic programming languages. In: Sutcliffe, G., Voronkov, A. (eds.) LPAR 2005. LNCS, vol. 3835, pp. 110–124. Springer, Heidelberg (2005)
20. Overton, D., Somogyi, Z., Stuckey, P.J.: Constraint-based mode analysis of Mercury. In: PPDP, pp. 109–120. ACM, New York (2002)
21. Pfenning, F.: *Types in logic programming*. MIT Press, Cambridge (1992)
22. Simons, P., Niemelä, I., Sooinen, T.: Extending and implementing the stable model semantics. *Artif. Intell.* 138(1-2), 181–234 (2002)

A Syntax and Semantics of DLV

In this section we provide the bare minimum of definitions of the syntax and semantics of DLV [16] (disjunctive Datalog) and we briefly describe its extension with functions [5].

A *term* is either a variable or a constant. An *atom* a is an expression $p(t_1, \dots, t_n)$, where p is a predicate of arity n and t_1, \dots, t_n are terms. A *classical literal* is an atom of the form a or $\neg a$, where \neg is the classical negation. A *literal* l has the form l or *not* l , where l is a classical literal and *not* represents the *negation as failure* [17]. A *disjunctive rule* r is a formula of the form

$$A_1 \vee \dots \vee A_n \leftarrow B_1, \dots, B_k, \text{not } B_{k+1}, \dots, \text{not } B_m$$

where $A_1, \dots, A_n, B_1, \dots, B_m$ are classical literals, $n \geq 0$ and $m \geq k \geq 0$. The disjunction $A_1 \vee \dots \vee A_n$ is the *head* of r , the conjunction $B_1, \dots, B_k, \text{not } B_{k+1}, \dots, \text{not } B_m$ is the *body* of r . We use the following notations: $HD(r) = \{A_1, \dots, A_n\}$ (the *head literals*), $BD^+(r) = \{B_1, \dots, B_k\}$ (the *positive body*), $BD^-(r) = \{B_{k+1}, \dots, B_m\}$ (the *negative body*). A rule with empty head (i.e., $n = 0$) is called a *constraint*, a rule with empty body (i.e., $k = m = 0$) is called a *fact*. A rule r is *safe* if any variable occurring in r also appears in $BD^+(r)$. A *DLV program* is a finite set of safe rules. We recall some basics of logic programming: The Herbrand Universe U_P of P is the set of all constants occurring in P (we assume $U_P \neq \emptyset$). The Herbrand Base B_P of P is the set of all the literals constructible from the predicate symbols of P and the constants of U_P . For every rule r of P , $grnd(r)$ denotes the set of rules obtained by applying all possible substitutions from the variables in r to elements of U_P ; $grnd(P)$ is the union of the sets $grnd(r)$, for every r of P . An *interpretation* I of P is a consistent subset of B_P , i.e. I does not contain pairs of classical literals of the form a and $\neg a$. Let P be a positive program (namely, for every rule r of P , $BD^-(r) = \emptyset$) and let $I \subseteq B_P$ be an interpretation.

- I is *closed under* P if, for every $r \in grnd(P)$, $BD^+(r) \subseteq I$ implies $HD(r) \cap I \neq \emptyset$.
- I is an *answer set* for P if I is minimal under set inclusion among all the interpretations closed under P .

The *reduct* or *Gelfond-Lifschitz* transformation of a ground program P with respect to an interpretation $I \subseteq B_P$ is the positive ground program P^I obtained from P by:

- deleting all rules $r \in P$ such that $BD^-(r) \cap I \neq \emptyset$;
- deleting the negative body from the remaining rules.

An *answer set* for a program P is an interpretation $I \subseteq B_P$ such that I is an answer set of $grnd(P)^I$.

Recently, DLV has been extended to DLV-Complex by introducing functions, sets and lists. This is obtained through the concept of *value invention* [5], which is based on the possibility of using externally defined predicates in the body of clauses. External predicates must be (well) moded, e.g. $\#_p(i, o)$ denotes a call to an external predicates

$p(c, X)$ with ground input c that will “invent” an output value for X . The semantics of DLV programs has to be extended in order to treat external predicates, which are interpreted by means of external oracles. A DLV-Complex program with external predicates is required to be safe and *VI-restricted* [5]. In this case, it is guaranteed that the grounding process will halt with a finite ground program whose answer sets are the expected ones. A function symbol $f(X)$ is introduced in DLV-Complex by associating with it a constructor predicate $\#f(o, i)$ and a destructor predicate $\#f(i, o)$. Given a value c , $\#f(v, c)$ invents a value v representing the ground term $f(c)$, while $\#f(v, X)$ reconstructs the value c .