

Debugging ASP Programs by Means of ASP*

Martin Brain¹, Martin Gebser², Jörg Pührer³, Torsten Schaub^{2,**},
Hans Tompits³, and Stefan Woltran³

¹ Department of Computer Science, University of Bath,
Bath, BA2 7AY, United Kingdom
mjb@cs.bath.ac.uk

² Institut für Informatik, Universität Potsdam,
August-Bebel-Straße 89, D-14482 Potsdam, Germany
{gebser, torsten}@cs.uni-potsdam.de

³ Institut für Informationssysteme, Technische Universität Wien,
Favoritenstraße 9–11, A-1040 Vienna, Austria
{puehrer, tompits, stefan}@kr.tuwien.ac.at

Abstract. Answer-set programming (ASP) has become an important paradigm for declarative problem solving in recent years. However, to further improve the usability of answer-set programs, the development of software-engineering tools is vital. In particular, the area of debugging provides a challenge in both theoretical and practical terms. This is due to the purely declarative nature of ASP that, on the one hand, calls for solver-independent methodologies and, on the other hand, does not directly apply to tracing techniques. In this paper, we propose a novel methodology, which rests within ASP itself, to sort out errors on the conceptual level. Our method makes use of *tagging*, where the program to be analyzed is rewritten using dedicated control atoms. This provides a flexible way to specify different types of debugging requests and a first step towards a dedicated (meta level) debugging language.

1 Introduction

Answer-set programming (ASP) has become a popular approach to declarative problem solving. The highly declarative semantics of the language decouples the problem specification from the computation of a solution. As a consequence, there is no general handle on the solving process whenever the output is in question. This deprives us of applying standard, procedural debugging methodologies and has led to a significant lack of methods and tools for debugging logic programs in ASP. However, the semantics itself allows for debugging methodologies that explain *why*, rather than *how*, a program is wrong. Another challenge is that the specification of a problem and its solutions are expressed in different fragments of the underlying language. While an encoding is usually posed in terms of predicate and variable symbols, a solution is free of variables and consists of ground atomic formulas.

We address this gap by proposing a novel debugging methodology that allows for debugging logic programs by appeal to ASP itself. To this end, we exploit and further

* This work was partially supported by the Austrian Science Fund (FWF) under project P18019.

** Affiliated with Simon Fraser University, Canada, and Griffith University, Australia.

extend the *tagging technique* introduced by Delgrande, Schaub, and Tompits [1] for compiling ordered logic programs into standard ones. The idea is to compile a program in focus (once) and to subsequently accomplish different types of debugging requests by appeal to special debugging modules using dedicated control atoms, called *tags*. Tags allow for controlling the formation of answer sets and reflect different properties (like the applicability status of a rule, for instance) and therefore can be used for manipulating the evaluation of the program (like the actual application of a rule).

The basic tagging technique is then used to conceive a (meta level) *debugging language* providing dedicated debugging statements. The idea here is to first translate a program into its tagged form and then to analyze it by means of debugging statements. More specifically, starting with a program Π over an alphabet \mathcal{A} , Π is translated into a tagged kernel program $\mathcal{T}_K[\Pi]$ over an extended alphabet \mathcal{A}^+ , and a debugging request Δ , formulated in the debugging language, is then compiled into a tagged program $\mathcal{D}[\Delta]$ over \mathcal{A}^+ . The debugging results are eventually read off the answer sets of the combined tagged program $\mathcal{T}_K[\Pi] \cup \mathcal{D}[\Delta]$. In this paper, we focus on the basic constituents of such a debugging language, confining ourselves to a detailed account of the tagging method.

Our approach has several advantageous distinct features: Firstly, it is based on meta-programming techniques that keep it within the realm of ASP. The dedicated debugging language offers an easy and modular way of specifying debugging requests. Notably, it allows the users to pose their requests with variables, which provides a tight connection to an encoding at hand. Secondly, the different debugging techniques are derived from semantic principles and relate to different characterizations of answer set formation. As a consequence, we can ascribe meaning to different debugging outcomes as well as the underlying compilation techniques. This is nicely demonstrated by our extrapolation techniques that allow for debugging incoherent logic programs. Finally, our approach has been implemented within the tool `SPOCK`, which is publicly available at [2].

Our approach is not meant to be universal. For one thing, it aims at exploring the limits of debugging within the realm of ASP. For another, it nicely complements the majority of existing approaches, all of which are external to ASP [3,4,5,6]. Most of them rely on graph-based characterizations, in the simplest case dependency graphs, and use specific algorithms for analyzing such graphs.

2 Background

Given an alphabet \mathcal{A} , a (*normal*) *logic program* is a finite set of rules of form

$$a \leftarrow b_1, \dots, b_m, \text{not } c_{m+1}, \dots, \text{not } c_n, \quad (1)$$

where $a, b_i, c_j \in \mathcal{A}$ are *atoms* for $0 \leq i \leq m \leq j \leq n$. A *literal* is an atom a or its negation $\text{not } a$. For a rule r of form (1), let $\text{head}(r) = a$ be the *head* of r and $\text{body}(r) = \{b_1, \dots, b_m, \text{not } c_{m+1}, \dots, \text{not } c_n\}$ be the *body* of r . Furthermore, we define $\text{body}^+(r) = \{b_1, \dots, b_m\}$ and $\text{body}^-(r) = \{c_{m+1}, \dots, c_n\}$. The set of atoms occurring in a program Π is given by $\text{At}(\Pi)$. For regrouping rules sharing the same head a , we use $\text{def}(a, \Pi) = \{r \in \Pi \mid \text{head}(r) = a\}$. For uniformity, we assume that any integrity constraint $\leftarrow \text{body}(r)$ is expressed as a rule $w \leftarrow \text{body}(r), \text{not } w$, where

w is a globally new atom. Moreover, we allow nested expressions of form *not not a*, where a is some atom, in the body of rules. Such rules are identified with normal rules in which *not not a* is replaced by *not a**, where a^* is a globally new atom, together with an additional rule $a^* \leftarrow \text{not } a$. We also take advantage of (singular) *choice rules* of form $\{a\} \leftarrow \text{body}(r)$ [7], which are an abbreviation for $a \leftarrow \text{body}(r)$, *not not a*.

We define answer sets following the approach of Lin and Zhao [8]. Given a program Π , let $PF(\Pi) \cup CF(\Pi)$ be the *completion* of Π [9], where

$$PF(\Pi) = \{ \text{body}(r) \rightarrow \text{head}(r) \mid r \in \Pi \} \text{ and}$$

$$CF(\Pi) = \{ a \rightarrow \bigvee_{r \in \text{def}(a, \Pi)} \text{body}(r) \mid a \in \mathcal{A} \}.$$
¹

A *loop* is a (non-empty) set of atoms that circularly depend upon each other in a program's positive atom dependency graph [8]. Programs having an acyclic positive atom dependency graph are *tight* [10]. The *loop formula* associated with a loop L is

$$LF(\Pi, L) = \neg(\bigvee_{r \in R(\Pi, L)} \text{body}(r)) \rightarrow \bigwedge_{a \in L} \neg a,$$

where $R(\Pi, L) = \{r \in \Pi \mid \text{head}(r) \in L, \text{body}^+(r) \cap L = \emptyset\}$. We denote the set of all loops in Π by $\text{loop}(\Pi)$. The set of all loop formulas of Π is $LF(\Pi) = \{LF(\Pi, L) \mid L \in \text{loop}(\Pi)\}$. A set X of atoms is an *answer set* of a logic program Π iff X is a model of $PF(\Pi) \cup CF(\Pi) \cup LF(\Pi)$. If Π is tight, then the answer sets of Π coincide with the models of $PF(\Pi) \cup CF(\Pi)$ (models of the latter are also referred to as the *supported models* of Π). The set Π_X of *generating rules* of a set X of atoms from program Π is defined as $\{r \in \Pi \mid \text{body}^+(r) \subseteq X, \text{body}^-(r) \cap X = \emptyset\}$.

As an example, consider $\Pi_1 = \{a \leftarrow; c \leftarrow \text{not } b, \text{not } d; d \leftarrow a, \text{not } c\}$ and its two answer sets $\{a, c\}$ and $\{a, d\}$. The completion of Π_1 is logically equivalent to $a \wedge \neg b \wedge (c \leftrightarrow \neg b \wedge \neg d) \wedge (d \leftrightarrow a \wedge \neg c)$; its models coincide with the answer sets of Π_1 . Adding $\{b \leftarrow e; e \leftarrow b\}$ to Π_1 induces the loop $\{b, e\}$ but leaves the set of answer sets of Π_1 intact. Unlike this, the completion of Π_1 becomes $a \wedge (b \leftrightarrow e) \wedge (c \leftrightarrow \neg b \wedge \neg d) \wedge (d \leftrightarrow a \wedge \neg c)$ and admits an additional model $\{a, b, d, e\}$. This supported model violates the loop formula $LF(\Pi_1, \{b, e\}) = \top \rightarrow \neg b \wedge \neg e$, which denies it the status of an answer set.

3 Debugging Modules

Our approach relies on the *tagging technique* introduced by Delgrande, Schaub, and Tompits [1] for compiling ordered logic programs back into normal programs. The idea is to rewrite a program by introducing so-called *tags* that allow for controlling the formation of answer sets. More formally, given a logic program Π over \mathcal{A} and a set \mathcal{N} of names for all rules in Π , we consider an enriched alphabet \mathcal{A}^+ obtained from \mathcal{A} by adding new pairwise distinct propositional atoms such as $\text{ap}(n_r)$, $\text{bl}(n_r)$, $\text{ok}(n_r)$, $\text{ko}(n_r)$, etc., where $n_r \in \mathcal{N}$ for each $r \in \Pi$. Intuitively, $\text{ap}(n_r)$ and $\text{bl}(n_r)$ express whether a rule r is applicable or blocked, respectively, while $\text{ok}(n_r)$ and $\text{ko}(n_r)$ are used for manipulating the application of r . Further tags are introduced in the sequel.

¹ Strictly speaking, $CF(\Pi)$ should take \mathcal{A} as additional argument; for simplicity, we leave this implicit. Moreover, $\text{body}(r)$ is understood as a conjunction of (classical) literals within $PF(\Pi)$, $CF(\Pi)$, and $LF(\Pi, L)$ in what follows.

3.1 Kernel Debugging Module

Our *kernel translation*, \mathcal{T}_K , decomposes rules of a given program such that they can be accessed by tags for controlling purposes.

Definition 1. *Let Π be a logic program over \mathcal{A} . Then, the program $\mathcal{T}_K[\Pi]$ over \mathcal{A}^+ consists of the following rules, for $r \in \Pi$, $b \in \text{body}^+(r)$, and $c \in \text{body}^-(r)$:*

$$\begin{aligned} \text{head}(r) &\leftarrow \text{ap}(n_r), \text{ not } \text{ko}(n_r), & \text{bl}(n_r) &\leftarrow \text{ok}(n_r), \text{ not } b, \\ \text{ap}(n_r) &\leftarrow \text{ok}(n_r), \text{ body}(r), & \text{bl}(n_r) &\leftarrow \text{ok}(n_r), \text{ not } \text{not } c, \\ & & \text{ok}(n_r) &\leftarrow \text{not } \overline{\text{ok}}(n_r). \end{aligned}$$

An auxiliary atom $\text{ap}(n_r)$, $\text{bl}(n_r)$, or $\text{ok}(n_r)$, respectively, occurs in an answer set X of $\mathcal{T}_K[\Pi]$ only if $r \in \Pi$. Also, for any $r \in \Pi$, there is a priori no atom $\text{ko}(n_r)$ contained in an answer set of $\mathcal{T}_K[\Pi]$, whereas $\text{ok}(n_r)$ is contained in any answer set of $\mathcal{T}_K[\Pi]$ by default. The role of $\overline{\text{ok}}(n_r)$ is to implement potential changes of this default behavior.

The following proposition collects more interesting relations.

Proposition 1. *Let Π be a logic program over \mathcal{A} and X an answer set of $\mathcal{T}_K[\Pi]$. Then, for any $r \in \Pi$ and $a \in \mathcal{A}$:*

1. $\text{ap}(n_r) \in X$ iff $r \in \Pi_X$ iff $\text{bl}(n_r) \notin X$;
2. if $a \in X$, then $\text{ap}(n_r) \in X$ for some $r \in \text{def}(a, \Pi)$;
3. if $a \notin X$, then $\text{bl}(n_r) \in X$ for all $r \in \text{def}(a, \Pi)$.

The relation between auxiliary atoms and original atoms from \mathcal{A} is described next.

Theorem 1. *Let Π be a logic program over \mathcal{A} . We have a one-to-one correspondence between the answer sets of Π and $\mathcal{T}_K[\Pi]$ satisfying the following conditions:*

1. If X is an answer set of Π , then

$$X \cup \{\text{ok}(n_r) \mid r \in \Pi\} \cup \{\text{ap}(n_r) \mid r \in \Pi_X\} \cup \{\text{bl}(n_r) \mid r \in \Pi \setminus \Pi_X\}$$

is an answer set of $\mathcal{T}_K[\Pi]$.

2. If Y is an answer set of $\mathcal{T}_K[\Pi]$, then $(Y \cap \mathcal{A})$ is an answer set of Π .

3.2 Extrapolating Non-existing Answer Sets

Whenever a program Π has no answer set, this means in terms of the characterization by Lin and Zhao [8] that there is no interpretation jointly satisfying $PF(\Pi)$, $CF(\Pi)$, and $LF(\Pi)$. In other words, each interpretation X over \mathcal{A} causes at least one of the following problems:

- If X falsifies $PF(\Pi)$, then there is some rule r in Π such that $\text{body}^+(r) \subseteq X$ and $\text{body}^-(r) \cap X = \emptyset$, but $\text{head}(r) \notin X$.
- If X falsifies $CF(\Pi)$, then there is some atom a in X that lacks a supporting rule, that is, $\text{body}^+(r) \not\subseteq X$ or $\text{body}^-(r) \cap X \neq \emptyset$ for each $r \in \text{def}(a, \Pi)$.
- If X falsifies $LF(\Pi)$, then X contains some loop L in Π that is unfounded with respect to X , that is, $X \not\models LF(\Pi, L)$.

This intuition is captured in the debugging model described below. It aims at analyzing incoherent situations by figuring out which rules or atoms cause some of the aforementioned problems. The names of the translations, \mathcal{T}_P , \mathcal{T}_C , and \mathcal{T}_L , reflect their respective purpose, indicating problems originating from the program, its completion, or its (non-trivial) loop formulas. We use abnormality atoms with a corresponding naming schema to indicate the respective problem: $ab_p(n_r)$ signals that rule r is falsified under some interpretation, $ab_c(a)$ points out that atom a is true but has no support, and $ab_l(a)$ aims at indicating an unfounded atom a .

Definition 2. *Let Π be a logic program over \mathcal{A} and A a set of atoms. Then:*

1. *The logic program $\mathcal{T}_P[\Pi]$ over \mathcal{A}^+ consists of the following rules, for all $r \in \Pi$:*

$$\begin{aligned} \{ \text{head}(r) \} \leftarrow \text{ap}(n_r), & \quad \text{ab}_p(n_r) \leftarrow \text{ap}(n_r), \text{not } \text{head}(r), \\ \text{ko}(n_r) \leftarrow . & \end{aligned}$$

2. *The logic program $\mathcal{T}_C[\Pi, A]$ over \mathcal{A}^+ consists of the following rules, for all $a \in A$, where $\{r_1, \dots, r_k\} = \text{def}(a, \Pi)$:*

$$\{ a \} \leftarrow \text{bl}(n_{r_1}), \dots, \text{bl}(n_{r_k}), \quad \text{ab}_c(a) \leftarrow a, \text{bl}(n_{r_1}), \dots, \text{bl}(n_{r_k}).$$

3. *The logic program $\mathcal{T}_L[A]$ over \mathcal{A}^+ consists of the following rules, for all $a \in A$:*

$$\{ \text{ab}_l(a) \} \leftarrow \text{not } \text{ab}_c(a), \quad a \leftarrow \text{ab}_l(a).$$

The purpose of adding facts $(\text{ko}(n_r) \leftarrow)$ in $\mathcal{T}_P[\Pi]$ is to avoid the application of the rule $(\text{head}(r) \leftarrow \text{ap}(n_r), \text{not } \text{ko}(n_r))$ in $\mathcal{T}_K[\Pi]$ (rather than to enforce a re-compilation in conjunction with $\mathcal{T}_K[\Pi]$). Regarding $\mathcal{T}_C[\Pi, A]$, note that $\text{def}(a, \Pi)$ might be empty, in which case we obtain for a the choice rule $(\{ a \} \leftarrow)$ along with $(\text{ab}_c(a) \leftarrow a)$. Observe that $\mathcal{T}_L[A]$ allows us to add $\text{ab}_l(a)$ to an answer set if a is supported. In contrast to $\text{ab}_p(n_r)$ in $\mathcal{T}_P[\Pi]$ and $\text{ab}_c(a)$ in $\mathcal{T}_C[\Pi, A]$, indicating violations of $PF(\Pi)$ or $CF(\Pi)$, respectively, the presence of $\text{ab}_l(a)$ in an answer set does not necessarily indicate the violation of any loop formula in $LF(\Pi)$. In fact, as the number of loops for Π can be exponential, we cannot reasonably check loop formula violations within $\mathcal{T}_L[A]$ (via an additional argument Π and tagged rules for analyzing loop formulas). Rather, as discussed below, we filter occurrences of $\text{ab}_l(a)$ in answer sets by *minimization*.

Next, we put things together. The answer sets of the subsequent translation are thought of as extrapolations of putative yet non-existing answer sets of the original program Π . That is, an atom $\text{ab}_p(n_r)$ signals that an answer set could be obtained if rule r was not contained in Π . Dually, $\text{ab}_c(a)$ indicates that an answer set could be obtained if atom a would be supported, that is, if it would be derivable by some rule. Finally, $\text{ab}_l(a)$ points to the violation of a loop formula that involves a . Moreover, we provide further possibilities to parametrize a debugging request by two additional arguments, Π' and A . Hereby, Π' restricts the set of rules (from program Π) whose violation is tolerated for debugging purposes, while A restricts the atoms that can be assumed true though being unsupported or belonging to a (non-trivial) unfounded set.

Definition 3. Let Π be a logic program over \mathcal{A} , $\Pi' \subseteq \Pi$, and $A \subseteq \text{At}(\Pi)$. Then:

$$\mathcal{T}_E[\Pi, \Pi', A] = \mathcal{T}_K[\Pi] \cup \mathcal{T}_P[\Pi'] \cup \mathcal{T}_C[\Pi, A] \cup \mathcal{T}_L[A].$$

Moreover, let $\mathcal{T}_E[\Pi, \Pi'] = \mathcal{T}_E[\Pi, \Pi', \text{At}(\Pi')]$ and $\mathcal{T}_E[\Pi] = \mathcal{T}_E[\Pi, \Pi, \text{At}(\Pi)]$.

Note that $\mathcal{T}_L[A]$ can be omitted in the definition of $\mathcal{T}_E[\Pi, \Pi', A]$ if Π is tight.

We list some basic properties first.

Proposition 2. Let Π be a logic program over \mathcal{A} and X an answer set of $\mathcal{T}_E[\Pi]$. Then, for each $r \in \Pi$:

1. $\text{ab}_p(n_r) \in X$ iff $\text{ap}(n_r) \in X$, $\text{bl}(n_r) \notin X$, and $\text{head}(r) \notin X$;
2. $\text{ab}_p(n_r) \notin X$ if $\text{ab}_c(\text{head}(r)) \in X$ or $\text{ab}_l(\text{head}(r)) \in X$.

Moreover, for every $a \in \text{At}(\Pi)$, it holds that:

1. $\text{ab}_c(a) \in X$ and $\text{ab}_l(a) \notin X$ iff $a \in X$ and $(X \cap \mathcal{A}) \not\models (\bigvee_{r \in \text{def}(a, \Pi)} \text{body}(r))$;
2. $\text{ab}_c(a) \notin X$ if $a \in X$ and $(X \cap \mathcal{A}) \models (\bigvee_{r \in \text{def}(a, \Pi)} \text{body}(r))$;
3. $\text{ab}_c(a) \notin X$ and $\text{ab}_l(a) \notin X$ if $a \notin X$;
4. $\text{ab}_c(a) \notin X$ if $\text{ab}_l(a) \in X$.

The next result shows that abnormality-free answer sets of the translated program correspond to the answer sets of the original program.² To this end, we introduce, for a program Π , $\text{AB}(\Pi) = (\{\text{ab}_p(n_r) \mid r \in \Pi\} \cup \{\text{ab}_c(a), \text{ab}_l(a) \mid a \in \text{At}(\Pi)\})$.

Theorem 2. Let Π be a logic program over \mathcal{A} . Then, it holds that:

1. If X is an answer set of Π , then

$$X \cup \{\text{ok}(n_r), \text{ko}(n_r) \mid r \in \Pi\} \cup \{\text{ap}(n_r) \mid r \in \Pi_X\} \cup \{\text{bl}(n_r) \mid r \in \Pi \setminus \Pi_X\}$$

is an answer set of $\mathcal{T}_E[\Pi]$.

2. If Y is an answer set of $\mathcal{T}_E[\Pi]$ such that $(Y \cap \text{AB}(\Pi)) = \emptyset$, then $(Y \cap \mathcal{A})$ is an answer set of Π .

The interesting case, however, is when the original program is incoherent. For illustrating this, let us consider three simple examples. To begin with, consider:

$$\Pi_2 = \{ n_{r_1} : a \leftarrow, n_{i_1} : \leftarrow a \}.$$

The program $\mathcal{T}_K[\Pi_2]$ consists of the following rules:

$$\begin{array}{ll} a \leftarrow \text{ap}(n_{r_1}), \text{not } \text{ko}(n_{r_1}), & \leftarrow \text{ap}(n_{i_1}), \text{not } \text{ko}(n_{i_1}), \\ \text{ap}(n_{r_1}) \leftarrow \text{ok}(n_{r_1}), & \text{ap}(n_{i_1}) \leftarrow \text{ok}(n_{i_1}), a, \\ & \text{bl}(n_{i_1}) \leftarrow \text{ok}(n_{i_1}), \text{not } a, \\ \text{ok}(n_{r_1}) \leftarrow \text{not } \overline{\text{ok}}(n_{r_1}), & \text{ok}(n_{i_1}) \leftarrow \text{not } \overline{\text{ok}}(n_{i_1}). \end{array}$$

² Due to relaxing Π by tolerating abnormalities, $\mathcal{T}_E[\Pi]$ always admits (abnormal) answer sets.

We obtain $\mathcal{T}_E[\Pi_2, \{r_1\}]$ by adding $(\mathcal{T}_P[\{r_1\}] \cup \mathcal{T}_C[\Pi_2, \{a\}] \cup \mathcal{T}_L[\{a\}])$, given next:

$$\begin{array}{ll} \text{ko}(n_{r_1}) \leftarrow , & \{a\} \leftarrow \text{bl}(n_{r_1}) , \\ & \text{ab}_c(a) \leftarrow a, \text{bl}(n_{r_1}) , \\ \{a\} \leftarrow \text{ap}(n_{r_1}) , & \{\text{ab}_l(a)\} \leftarrow \text{not } \text{ab}_c(a) , \\ \text{ab}_p(n_{r_1}) \leftarrow \text{ap}(n_{r_1}), \text{not } a , & a \leftarrow \text{ab}_l(a) . \end{array}$$

The unique answer set of $\mathcal{T}_E[\Pi_2, \{r_1\}]$ is:³

$$\{\underline{\text{ab}}_p(n_{r_1}), \text{ap}(n_{r_1}), \text{bl}(n_{i_1}), \text{ok}(n_{r_1}), \text{ok}(n_{i_1}), \text{ko}(n_{r_1})\} . \quad (2)$$

Note that applying the modules from Definition 2 only to subprogram $\{r_1\}$ makes us focus on answer set candidates satisfying the residual program $\{i_1\}$. The abnormality tag $\text{ab}_p(n_{r_1})$ signals that, in order to obtain an answer set, rule r_1 must not be applied.

As another example, consider:

$$\Pi_3 = \{ n_{r_1} : a \leftarrow b, n_{i_1} : \leftarrow \text{not } a \} .$$

Program $\mathcal{T}_E[\Pi_3, \{r_1\}, \{a\}]$ has the unique answer set:

$$\{a, \underline{\text{ab}}_c(a), \text{bl}(n_{r_1}), \text{bl}(n_{i_1}), \text{ok}(n_{r_1}), \text{ok}(n_{i_1}), \text{ko}(n_{r_1})\} , \quad (3)$$

indicating that a lacks supporting rules.

Finally, consider the following program:

$$\Pi_4 = \{ n_{r_1} : a \leftarrow b, n_{r_2} : b \leftarrow a, n_{i_1} : \leftarrow \text{not } a \} .$$

Program $\mathcal{T}_E[\Pi_4, \{r_1, r_2\}]$ has four answer sets (omitting $\text{ok}(n_{r_1}), \text{ok}(n_{r_2}), \text{ok}(n_{i_1})$):

$$\{a, \underline{\text{ab}}_p(n_{r_2}), \underline{\text{ab}}_c(a), \text{bl}(n_{r_1}), \text{ap}(n_{r_2}), \text{bl}(n_{i_1}), \text{ko}(n_{r_1}), \text{ko}(n_{r_2})\} , \quad (4)$$

$$\{a, b, \underline{\text{ab}}_l(a), \text{ap}(n_{r_1}), \text{ap}(n_{r_2}), \text{bl}(n_{i_1}), \text{ko}(n_{r_1}), \text{ko}(n_{r_2})\} , \quad (5)$$

$$\{a, b, \underline{\text{ab}}_l(b), \text{ap}(n_{r_1}), \text{ap}(n_{r_2}), \text{bl}(n_{i_1}), \text{ko}(n_{r_1}), \text{ko}(n_{r_2})\} , \quad (6)$$

$$\{a, b, \underline{\text{ab}}_l(a), \underline{\text{ab}}_l(b), \text{ap}(n_{r_1}), \text{ap}(n_{r_2}), \text{bl}(n_{i_1}), \text{ko}(n_{r_1}), \text{ko}(n_{r_2})\} . \quad (7)$$

The last example illustrates that the relaxation mechanisms underlying translation \mathcal{T}_E can lead to overly involved explanations of the source of incoherence, as manifested by the first and last answer set (cf. (4) and (7)). Therefore, we suggest focusing on answer sets containing a *minimum number* of instances of ab predicates.⁴ In the last case, this gives the second and third answer set (cf. (5) and (6)). Indeed, adding only one fact, either $(a \leftarrow)$ or $(b \leftarrow)$, to the (untagged) incoherent program Π_4 makes it coherent.

The next results shed some more light on the semantic links between the original and the transformed program, whenever the former admits no answer set.

³ In what follows, we underline abnormality tags.

⁴ This can be implemented via `minimize` statements as available in `Smodels`.

Theorem 3. *Let Π be a logic program over \mathcal{A} . Then, it holds that:*

1. *If Y is an answer set of $\mathcal{T}_E[\Pi]$ and $\text{ab}_p(n_r) \in Y$, then $(Y \cap \mathcal{A}) \not\models (\text{body}(r) \rightarrow \text{head}(r))$, where $(\text{body}(r) \rightarrow \text{head}(r)) \in \text{PF}(\Pi)$;*
2. *If Y is an answer set of $\mathcal{T}_E[\Pi]$ and $\text{ab}_c(a) \in Y$, then $(Y \cap \mathcal{A}) \not\models (a \rightarrow \bigvee_{r \in \text{def}(a, \Pi)} \text{body}(r))$, where $(a \rightarrow \bigvee_{r \in \text{def}(a, \Pi)} \text{body}(r)) \in \text{CF}(\Pi)$;*
3. *If Y is an answer set of $\mathcal{T}_E[\Pi]$ such that, for some $L \in \text{loop}(\Pi)$, we have $L \subseteq (Y \cap \mathcal{A})$, $(Y \cap \mathcal{A}) \not\models \text{LF}(\Pi, L)$, and $(Y \cap \mathcal{A}) \models (\bigvee_{r \in \text{def}(a, \Pi)} \text{body}(r))$ for every $a \in L$, then $\text{ab}_l(a') \in Y$ for some $a' \in L$.*

The same results hold for partial compilations, $\mathcal{T}_E[\Pi, \Pi', A]$, but are omitted for brevity.

For illustration, let us return to the last three examples. Intersecting the only answer set (2) of $\mathcal{T}_E[\Pi_2, \{r_1\}]$ with the alphabet of Π_2 yields the empty set. We obtain $\emptyset \not\models (a \leftarrow)$, as indicated by $\text{ab}_p(n_{r_1})$ in (2). Note that the empty set is the only subset of $\text{At}(\Pi_2)$ that satisfies integrity constraint $i_1 \in \Pi_2$. Proceeding analogously with the only answer set (3) of $\mathcal{T}_E[\Pi_3, \{r_1\}, \{a\}]$ yields $\{a\}$, and we obtain $\{a\} \not\models (a \rightarrow b)$, as signaled by $\text{ab}_c(a)$ in (3). In fact, $\{a\}$ is the only subset of the atoms subject to extrapolation that satisfies integrity constraint $i_1 \in \Pi_3$.

Finally, consider the two abnormality-minimum answer sets (5) and (6) of $\mathcal{T}_E[\Pi_4, \{r_1, r_2\}]$. Both are actually symmetrical since they refer to the same loop $\{a, b\}$ through different elements, as indicated by $\text{ab}_l(a)$ and $\text{ab}_l(b)$, respectively. Hence, both answer sets (5) and (6) of $\mathcal{T}_E[\Pi_4, \{r_1, r_2\}]$ induce candidate set $\{a, b\}$, which falsifies its own loop formula: $\{a, b\} \not\models (\top \rightarrow \neg a \wedge \neg b)$. Note that $\{a\}$ is actually another candidate subset of $\text{At}(\Pi_4)$. However, this candidate suffers from two abnormalities, as indicated by the non-minimum answer set (4) through $\text{ab}_p(n_{r_2})$ and $\text{ab}_c(a)$. In fact, we have $\{a\} \not\models (b \leftarrow a)$, violating r_2 , and $\{a\} \not\models (a \rightarrow b)$, violating the completion.

The next result captures a more realistic scenario, in which only a subset of a program is subject to extrapolation and only abnormality-minimum answer sets of the translation are considered. From the perspective of an original program Π , the abnormality-minimum answer sets of $\mathcal{T}_E[\Pi, \Pi']$ provide us with the candidate sets among $\text{At}(\Pi)$ that satisfy the requirement of being an answer set of Π under a minimum number of repairs on Π' . A repair is either the deletion of a rule r or an addition of a fact $(a \leftarrow)$ (which prevents a from being not supported or unfounded). The former repair refers to $\text{ab}_p(n_r)$, which is clearly avoided when r is deleted, and the latter to $\text{ab}_c(a)$ or $\text{ab}_l(a)$, since $(a \rightarrow \bigvee_{r \in \text{def}(a, \Pi \cup \{a \leftarrow\})} \text{body}(r))$ and $\text{LF}(\Pi \cup \{a \leftarrow\}, L)$, for any loop L containing a , then amount to $(a \rightarrow \top)$ and $(\perp \rightarrow \bigwedge_{a \in L} \neg a)$.

Theorem 4. *Let Π be a logic program over \mathcal{A} and (Π_1, Π_2) a partition of Π such that $(\{\text{head}(r_1) \mid r_1 \in \Pi_1\} \cap \text{At}(\Pi_2)) = \emptyset$. Furthermore, let \mathcal{M} be the set of all answer sets Y of $\mathcal{T}_E[\Pi, \Pi_2]$ such that the cardinality of $(Y \cap \text{AB}(\Pi_2))$ is minimum among all answer sets of $\mathcal{T}_E[\Pi, \Pi_2]$. Then, it holds that:*

1. *If $Y \in \mathcal{M}$, then $(Y \cap \mathcal{A})$ satisfies all formulas in $(\text{PF}(\Pi_1) \cup (\text{CF}(\Pi_1) \setminus \{a \rightarrow \perp \mid a \in \text{At}(\Pi_2)\})) \cup \text{LF}(\Pi_1)$ and all formulas in $(\text{PF}(\Pi_2) \cup \text{CF}(\Pi_2) \cup \text{LF}(\Pi_2))$ under a minimum number of repairs on Π_2 ;*
2. *If $X \subseteq \mathcal{A}$ satisfies all formulas in $(\text{PF}(\Pi_1) \cup (\text{CF}(\Pi_1) \setminus \{a \rightarrow \perp \mid a \in \text{At}(\Pi_2)\})) \cup \text{LF}(\Pi_1)$ and all formulas in $(\text{PF}(\Pi_2) \cup \text{CF}(\Pi_2) \cup \text{LF}(\Pi_2))$ under a minimum number of repairs on Π_2 , then there is a $Y \in \mathcal{M}$ such that $X = (Y \cap \mathcal{A})$.*

3.3 Ordering Rule Applications

Our last debugging component allows for imposing an order on the rule applications. To this end, we make use of ordered logic programs following the framework of Delgrande, Schaub, and Tompits [1].

Definition 4. Let Π be a logic program over \mathcal{A} and let $<$ be a strict partial order over Π . Then, $\mathcal{T}_o[\Pi, <]$ consists of the following rules, for all $r, r' \in \Pi$ with $r < r'$ and $\{r_1, \dots, r_k\} = \{r'' \mid r < r''\}$:

$$\begin{aligned} \overline{\text{ok}}(n_r) &\leftarrow , & \text{rdy}(n_r, n_{r'}) &\leftarrow \text{ap}(n_{r'}) , \\ \text{ok}(n_r) &\leftarrow \text{rdy}(n_r, n_{r_1}), \dots, \text{rdy}(n_r, n_{r_k}) , & \text{rdy}(n_r, n_{r'}) &\leftarrow \text{bl}(n_{r'}) . \end{aligned}$$

Furthermore, $\mathcal{T}_o[<]$ is a shortcut for $\mathcal{T}_o[\Pi_{<}, <]$, where $\Pi_{<} = \{r, r' \in \Pi \mid r < r'\}$.

The order $<$ on Π singles out the answer sets of Π whose rule application and/or blockage is compatible with $<$. That is, given the order $r_1 < r_2$, the higher-ranked rule r_2 must be applied or found to be blocked before r_1 . In other words, $<$ selects those answer sets X of Π that can be generated by an appropriate sequence of elements of Π_X . We refer for a detailed formal elaboration to Delgrande, Schaub, and Tompits [1].⁵

3.4 Debugging Programs with Variables

With two exceptions, the translations discussed so far carry over to programs with variables simply by using parametrized names. To this end, replace every occurrence of a name n_r by $n_r(\mathbf{X})$, where \mathbf{X} is the sequence of variables occurring in rule r . Accordingly, a rule $(r : p(X, Y) \leftarrow q(1, X), \text{not } s(Y))$ gets the name $n_r(X, Y)$, for instance, yielding

$$\begin{aligned} \text{ap}(n_r(X, Y)) &\leftarrow \text{ok}(n_r(X, Y)), q(1, X), \text{not } s(Y) & \text{and} \\ \{p(X, Y)\} &\leftarrow \text{ap}(n_r(X, Y)) . \end{aligned}$$

Similarly, we get for an atomic formula $p(X, Y)$:

$$\{\text{ab}_l(p(X, Y))\} \leftarrow \text{not } \text{ab}_c(p(X, Y)) .$$

The final representation (e.g., with or without function symbols) of an atomic formula like $\text{ab}_l(p(X, Y))$ depends on the language capacities of the target ASP solver.

The first exception is resolved by replacing the rules in Item 2 of Definition 2 by

$$\{a(\mathbf{X})\} \leftarrow \text{not } \text{ap}_x(a(\mathbf{X})) \quad \text{and} \quad \text{ab}_c(a(\mathbf{X})) \leftarrow a(\mathbf{X}), \text{not } \text{ap}_x(a(\mathbf{X})) ,$$

and by adding, for each $r \in \text{def}(a, \Pi)$, the rule

$$\text{ap}_x(a(\mathbf{X})) \leftarrow \text{ap}(n_r(\mathbf{Y})) ,$$

⁵ For obtaining the precise semantics of [1], predicate bl must actually be replaced by another predicate bl^* using no nested expressions like $\text{not not } c$. The definition of bl^* is omitted here.

where apx is an auxiliary predicate symbol and \mathbf{X} is a subsequence of the variables in \mathbf{Y} . An atom $\text{apx}(a(\mathbf{c}))$ indicates that $a(\mathbf{c})$ is not derivable given that all putative rules in $\text{def}(a, \Pi)$ are inapplicable.

The second exception concerns translation \mathcal{T}_o which necessitates that the set of dominating rules $\{r'' \mid r < r''\}$ is ground in order to guarantee that all instances have been applied or blocked (cf. Definition 4). Unlike this, the name of r , n_r , can be parametrized.

Apart from firmer instantiations, further optimizations are possible by use of domain predicates for names and atoms, like *name* and *atom*. For instance, a rule $(\text{ok}(n_r) \leftarrow \text{not } \overline{\text{ok}}(n_r))$ could be represented as $(\text{ok}(X) \leftarrow \text{name}(X), \text{not } \overline{\text{ok}}(X))$.

3.5 Implementation

The tool `spock` implements a collection of methods for debugging ASP programs, following the ideas introduced above. The system can be used either with DLV [11] or with `Smodels` [7] (together with `Lparse`) and is obtainable at [2], where also some further information about the system is available.

The tool is written in Java 5.0 and published under the GNU General Public License. The input of `spock` is a logic program in the core language of either DLV or `Smodels`, read from the system standard input and/or from (multiple) files. The output varies according to the selected functionalities, determined by a set of options. The most important syntax extension of the input programs, however, is the labeling of rules, allowing debugging mechanisms to explicitly refer to certain rules.

4 Elements of a Debugging Language

In order to enhance the usability and convenience of our debugging technique, we provide in this section some basic elements for a higher-level debugging language. The idea is to compile a program once and to explore it subsequently by means of debugging statements. To this end, we assume that the program Π has been compiled into $\mathcal{T}_K[\Pi]$. A debugging request, Δ , formulated in the debugging language, is compiled by means of a function \mathcal{D} into a tagged program $\mathcal{D}[\Delta]$ such that the debugging results are read off the answer sets of the combined tagged program $\mathcal{T}_K[\Pi] \cup \mathcal{D}[\Delta]$. We stress that the subsequent discussion is intended as a starting point only, given at a rather informal level; a detailed elaboration of the full language will be explored elsewhere.

The first type of expressions are referred to as *enforcement statements* since they may alter the original set of answer sets. The following expressions are enforcement statements (in what follows, let C be a set of literals over $\text{At}(\Pi)$ and n_r the name of some $r \in \Pi$):

- **block** n_r **if** C ;
- **apply** n_r **if** C ;
- **assign** $a_1 = v_1, \dots, a_k = v_k$ **if** C , where $a_i \in \text{At}(\Pi)$ and $v_i \in \{\mathbf{t}, \mathbf{f}\}$ for $i = 1, \dots, k$.

The semantics of enforcement statements is given via a compilation function \mathcal{D} :

- $\mathcal{D}[\mathbf{block } n_r \text{ if } C] = \{\mathbf{ko}(n_r) \leftarrow C, \mathbf{bl}(n_r) \leftarrow C, \overline{\mathbf{ok}}(n_r) \leftarrow C\}$;
- $\mathcal{D}[\mathbf{apply } n_r \text{ if } C] = \{\mathbf{ap}(n_r) \leftarrow C, \overline{\mathbf{ok}}(n_r) \leftarrow C\}$;
- $\mathcal{D}[\mathbf{assign } a_1 = v_1, \dots, a_k = v_k \text{ if } C] = \{a_i \leftarrow C \mid v_i = \mathbf{t}\} \cup \{\leftarrow a_i, C \mid v_i = \mathbf{f}\}$.

For example, $\mathbf{block } n_r(X) \text{ if } \mathit{not } (X > 5)$ blocks rule r unless instantiated with objects larger than 5. Also, one may use tags within the precondition, as in $\mathbf{block } n_r(X) \text{ if } \mathbf{ap}(n_s(X))$.

Unlike the above, *projection statements* do not alter the original set of answer sets but return a specific subset of the original answer sets:

- $\mathbf{blocked } n_1, \dots, n_k \text{ if } C$;
- $\mathbf{applied } n_1, \dots, n_k \text{ if } C$;
- $\mathbf{assigned } a_1 = v_1, \dots, a_k = v_k \text{ if } C$, where $a_i \in \mathit{At}(II)$ and $v_i \in \{\mathbf{t}, \mathbf{f}\}$ for $i = 1, \dots, k$.

The semantics of projection statements is again given in terms of a compilation:

- $\mathcal{D}[\mathbf{blocked } n_1, \dots, n_k \text{ if } C] = \{\leftarrow \mathit{not } \mathbf{bl}(n_1), \dots, \mathit{not } \mathbf{bl}(n_k), C\}$;
- $\mathcal{D}[\mathbf{applied } n_1, \dots, n_k \text{ if } C] = \{\leftarrow \mathit{not } \mathbf{ap}(n_1), \dots, \mathit{not } \mathbf{ap}(n_k), C\}$;
- $\mathcal{D}[\mathbf{assigned } a_1 = v_1, \dots, a_k = v_k \text{ if } C] = \{\leftarrow \mathit{not } a_i, C \mid v_i = \mathbf{t}\} \cup \{\leftarrow a_i, C \mid v_i = \mathbf{f}\}$.

For example, $\mathbf{blocked } n_r(X), n_s(Y) \text{ if } X \neq Y$ eliminates answer sets to which rules r and s contribute with different instantiations. Applying the compilation function \mathcal{D} , we get

$$\mathcal{D}[\mathbf{blocked } n_r(X), n_s(Y) \text{ if } X \neq Y] = \{\leftarrow \mathit{not } \mathbf{bl}(n_r(X)), \mathit{not } \mathbf{bl}(n_s(Y)), X \neq Y\}.$$

Observe that the translation of the corresponding enforcement statement would yield multiple rules whose instantiations cannot be controlled in a synchronous fashion.

Next, we introduce expressions for analyzing incoherent situations by extrapolation, correspondingly called *extrapolation statements*:

- $\mathbf{extrapolate } n_r \text{ if } C$;
- $\mathbf{extrapolate}_x s \text{ if } C$, where $x \in \{\mathbf{p}, \mathbf{c}, \mathbf{l}\}$ and $s \in (\{n_r \mid r \in II\} \cup \mathit{At}(II))$;
- $\mathbf{minimize } X$, where $X \subseteq \{p, c, l\}$.

Further constructs, for instance, statements applying to entire sets of rules and/or atoms, are definable in a straightforward way but omitted for space reasons.

For a program II , the function \mathcal{D} defines the semantics of extrapolation statements:

- $\mathcal{D}[\mathbf{extrapolate } n_r \text{ if } C] = \mathcal{T}_E[II, \{r\}]$;
- $\mathcal{D}[\mathbf{extrapolate}_p s \text{ if } C] = \mathcal{T}_P[\{r\}]$, provided that $s \in \{n_r \mid r \in II\}$;
- $\mathcal{D}[\mathbf{extrapolate}_c s \text{ if } C] = \mathcal{T}_C[II, \{s\}]$, provided that $s \in \mathit{At}(II)$;
- $\mathcal{D}[\mathbf{extrapolate}_l s \text{ if } C] = \mathcal{T}_L[\{s\}]$, provided that $s \in \mathit{At}(II)$.

The semantics of the $\mathbf{minimize}$ command depends on the capacities of the underlying ASP solver. For instance, in `Smodels`, a global minimization of abnormalities, using a binary predicate `ab/2`, could be expressed as $\mathbf{minimize}\{\mathbf{ab}(X, Y)\}$.

A pragmatic variant of extrapolation is *variation*:

- **vary** a_1, \dots, a_k **if** C , where $a_i \in At(\Pi)$ for $i = 1, \dots, k$.

The semantics of variation statements is as follows:

- $\mathcal{D}[\mathbf{vary} a_1, \dots, a_k \mathbf{if} C] = \{ \{a_i\} \leftarrow C \mid i = 1, \dots, k \} \cup \{ \text{ko}(n_r) \leftarrow C \mid r \in \Pi, \text{head}(r) = a_i, i = 1, \dots, k \}$.

Finally, we introduce some procedural flavor, which allows for imposing a certain order of rule and/or atom considerations. A basic *ordinal statement* is an expression of the following form:

- s **before** t , where $s, t \in \{n_r \mid r \in \Pi\}$ or $s, t \in At(\Pi)$.

Its semantics is defined in the following way:

- $\mathcal{D}[s \mathbf{before} t] = \mathcal{T}_o[\{r_t < r_s\}]$, provided that $s = n_{r_s}$ and $t = n_{r_t}$;
- $\mathcal{D}[s \mathbf{before} t] = \mathcal{T}_o[\{r_t < r_s \mid \text{head}(r_s) = s, \text{head}(r_t) = t\}]$, provided that $s, t \in At(\Pi)$.

The (atom-based) generalization to sets is defined as A **before** B where $A, B \subseteq At(\Pi)$, and it could be used to debug a generate-and-test encoding by specifying the generate atoms' precedence over the test atoms.

5 Discussion and Related Work

Although debugging is mentioned as a possible application in many papers on program analysis (often with the implicit assertion that incoherent and erroneous programs are the same thing), there are relatively few papers that are primarily focused on debugging.

The `noMore` system [3] uses a graph-based algorithm for computing the answer sets of a program. The interface to the system allows the computation process to be visualized and animated, so that the user can observe certain parts of the computation process. Given some background knowledge on how answer-set computation algorithms work, this is an intuitive and appealing approach to debugging. However, as it works on ground programs, dealing with larger programs is difficult. Also, at the conceptual level, it blurs the distinction between what a program means and how its answer sets are computed.

Brain and De Vos [4] start by characterizing bugs as mismatches between what the programmer expects and the actual answer sets of a program. Two query algorithms, answering why some set S is in some answer set A and why some set S is not contained in any answer set, can then be used to explore these mismatches. The algorithms suggested are procedural and similar to the ones used in ASP solvers, suggesting that an approach using a program-level transformation would be more practical.

Syrjänen [5] discusses a pragmatic approach to debugging, focusing on incoherent ground programs, whose source of incoherence is an active constraint (or an odd negative cycle). Derived from the field of symbolic diagnosis [12], the corresponding system uses an approach similar to ours, using program transformations to find minimal sets of constraints that would be active and to find odd cycles.

Pontelli and Son [6] adopt the concept of *justifications* [13,14,15] to the context of ASP. Roughly, justifications formalize reasons why an atom is in an answer set.

One important area that is not considered by the existing approaches is that of interfaces to debugging systems. Most of the proposed methods produce very large amounts of structured information, and it is difficult to automatically identify which parts of it are of interest to the programmer [4]. Thus, the design of the debugging interface is critical to the utility of the finished system. It must allow the programmers to quickly and easily focus in on areas that they consider to be erroneous without overloading them with information. This is an open and important area of research if ASP is to achieve truly wide-spread use. Other directions of future work include extending the results given here to handle constructs such as cardinality rules, disjunction, aggregates, and functions. Also, there is a potential to use similar tagging systems to allow programs to reason about their own consistency. We mention *consistency restoring rules* [16] as one such example.

References

1. Delgrande, J., Schaub, T., Tompits, H.: A framework for compiling preferences in logic programs. *Theory and Practice of Logic Programming* **3**(2) (2003) 129–187
2. (<http://www.kr.tuwien.ac.at/research/debug>)
3. Bösel, A., Linke, T., Schaub, T.: Profiling answer set programming: The visualization component of the noMore system. In *Proc. JELIA'04*. Springer (2004) 702–705
4. Brain, M., De Vos, M.: Debugging logic programs under the answer set semantics. In *Proc. ASP'05*. (2005) 141–152
5. Syrjänen, T.: Debugging inconsistent answer set programs. In *Proc. NMR'06*. (2006) 77–83
6. Pontelli, E., Son, T.: Justifications for logic programs under answer set semantics. In *Proc. ICLP'06*. Springer (2006) 196–210
7. Simons, P., Niemelä, I., Sooinen, T.: Extending and implementing the stable model semantics. *Artificial Intelligence* **138**(1-2) (2002) 181–234
8. Lin, F., Zhao, Y.: ASSAT: Computing answer sets of a logic program by SAT solvers. *Artificial Intelligence* **157**(1-2) (2004) 115–137
9. Clark, K.: Negation as failure. In *Logic and Data Bases*. Plenum (1978) 293–322
10. Fages, F.: Consistency of Clark's completion and the existence of stable models. *Journal of Methods of Logic in Computer Science* **1** (1994) 51–60
11. Leone, N., Pfeifer, G., Faber, W., Eiter, T., Gottlob, G., Perri, S., Scarcello, F.: The DLV system for knowledge representation and reasoning. *ACM Transactions on Computational Logic* **7**(3) (2006) 499–562
12. Reiter, R.: A theory of diagnosis from first principles. *Artificial Intelligence* **32**(1) (1987) 57–95
13. Roychoudhury, A., Ramakrishnan, C., Ramakrishnan, I.: Justifying proofs using memo tables. In *Proc. PPDP'00*. (2000) 178–189
14. Pemmasani, G., Guo, H., Dong, Y., Ramakrishnan, C., Ramakrishnan, I.: Online justification for tabled logic programs. In *Proc. FLOPS'04*. Springer (2004) 24–38
15. Specht, G.: Generating explanation trees even for negations in deductive database systems. In *Proc. LPE'93*. (1993) 8–13
16. Balduccini, M., Gelfond, M.: Logic programs with consistency-restoring rules. In *Proc. Commonsense'03*. (2003) 9–18