


**I N F S Y S
R E S E A R C H
R E P O R T**



**INSTITUT FÜR INFORMATIONSSYSTEME
ARBEITSBEREICH WISSENSBASIERTE SYSTEME**

**IMPROVING HEX-PROGRAM EVALUATION
BASED ON UNFOUNDED SETS**

**THOMAS EITER MICHAEL FINK
THOMAS KRENNWALLNER CHRISTOPH REDL
PETER SCHÜLLER**

INFSYS RESEARCH REPORT DRAFT 1843-12-08

SEPTEMBER 2, 2013

Institut für Informationssysteme
AB Wissensbasierte Systeme
Technische Universität Wien
Favoritenstraße 9-11
A-1040 Wien, Austria
Tel: +43-1-58801-18405
Fax: +43-1-58801-18493
sek@kr.tuwien.ac.at
www.kr.tuwien.ac.at



kbs 
*Knowledge-Based
Systems Group*

IMPROVING HEX-PROGRAM EVALUATION BASED ON UNFOUNDED
SETSThomas Eiter¹ Michael Fink¹ Thomas Krennwallner¹ Christoph Redl¹
Peter Schüller²

Abstract. HEX programs extend logic programs with external computations through external atoms. As already reasoning from Horn programs with nonmonotonic external atoms of polynomial complexity is on the second level of the polynomial hierarchy, minimality checking of answer set candidates needs special attention. To this end, we present an approach based on *unfounded sets* as a generalization of related techniques for ordinary ASP programs. The unfounded set detection is efficiently expressed as a propositional SAT problem, for which we provide two different encodings and optimizations to them. We then integrate our approach into a previously developed evaluation framework for HEX-programs, which is enriched by additional learning techniques that aim at avoiding the reconstruction of the same or related unfounded sets. Furthermore, we provide a syntactic criterion that allows one to skip the minimality check in practically relevant cases. It benignly carries over to the components of an appropriate program decomposition, which can be exploited to restrict unfounded set search to relevant parts of the program. An experimental evaluation shows that the new approach significantly decreases runtime.

¹Institut für Informationssysteme, Technische Universität Wien, Favoritenstraße 9-11, A-1040 Vienna, Austria; email: {eiter,fink,tkren,redl}@kr.tuwien.ac.at.

²Faculty of Engineering and Natural Sciences, Sabanci University, Orhanli, Tuzla, 34956 Istanbul, Turkey; email: peterschueller@sabanciuniv.edu.

Acknowledgements: This research has been supported by the Vienna Science and Technology Fund (WWTF) project ICT08-020. This report is an extended version of [Eiter et al., 2012c,b,d].

Copyright © 2013 by the authors

Contents

1	Introduction	3
2	Preliminaries	5
2.1	HEX Programs	5
2.2	Evaluation	7
3	Unfounded Set Detection	8
3.1	Basic Encoding of the Unfounded Set Search	9
3.2	Uniform Encoding of the Unfounded Set Search	11
4	Optimization and Learning	14
4.1	Optimization	14
4.2	Learning Nogoods from Unfounded Sets	17
4.2.1	UFS-based Learning	17
4.2.2	Reduct-based Learning	18
5	Deciding the Necessity of the Minimality Check	18
5.1	Basic Decision Criterion	19
5.2	Program Decomposition	21
6	Implementation and Evaluation	22
6.1	Detailed Benchmark Description	22
6.2	Summary	29
7	Discussion	29
7.1	Related Work	29
7.2	Extensions	30
8	Conclusion	31
A	Benchmark Encodings	36
A.1	Abstract Argumentation	36
A.2	Conformant Planning	37
B	Proofs	39

1 Introduction

Answer Set Programming (ASP) is a declarative problem solving approach which due to expressive and efficient systems like SMODELs [Simons et al., 2002], dlv [Leone et al., 2006] and CLASP [Gebser et al., 2012], has been gaining popularity for many applications [Brewka et al., 2011]. However, current trends in computing, such as context awareness or distributed systems, raised the need for access to external sources in a program, which for instance on the Web ranges from light-weight data access (e.g., XML, RDF, or data bases) to knowledge-intensive formalisms (e.g., OWL ontologies).

To cater for this need, HEX-programs [Eiter et al., 2005] extend ASP with so-called external atoms, through which the user can couple any external information source with a logic program. Roughly, such atoms pass information from the program, given by predicate extensions, into an external source which returns output values of an (abstract) function that it computes. This convenient extension has been exploited for many applications, including querying data and ontologies, e.g. [Eiter et al., 2008b, Hoehndorf et al., 2007, Marano et al., 2010], e-government [Zirtiloğlu and Yolum, 2008], fuzzy answer set programming [Nieuwenborgh et al., 2007a], multi-context reasoning [Brewka and Eiter, 2007, Eiter et al., 2012e], user interface adaptation [Zakraoui and Zagler, 2011], or reasoning about actions and planning [Nieuwenborgh et al., 2007b, Basol et al., 2010], to mention some of them. It is highly expressive as recursive data exchange between the rules and external sources is possible.

The semantics of HEX-programs is model-based and formally defined in terms of answer sets using the FLP reduct [Faber et al., 2011]: an interpretation \mathbf{A} is an answer set of a program Π , if it is a subset-minimal model of the FLP-reduct $f\Pi^{\mathbf{A}}$ of Π wrt. \mathbf{A} , where $f\Pi^{\mathbf{A}}$ is the set of all rules whose body is satisfied by \mathbf{A} . For ordinary programs (without external atoms), this definition is equivalent to the one using the seminal GL-reduct in [Gelfond and Lifschitz, 1991], but has advantages for extensions with nonmonotonic aggregates [Faber et al., 2011] and the more general external atoms in HEX-programs.

Traditionally, a HEX-program Π is evaluated in two steps as follows. In the first step, external atoms are viewed as ordinary atoms (called *replacement atoms*) whose truth values are guessed by choice rules. The respective ordinary ASP program $\hat{\Pi}$ is evaluated using an ordinary ASP solver, and each answer set $\hat{\mathbf{A}}$ obtained is checked against the external sources, i.e., the guess is verified. In case of success \mathbf{A} is a model of the reduct $f\Pi^{\mathbf{A}}$, but it is yet unknown whether it is subset-minimal. This is ensured in the second step by an *FLP check*. A straightforward method, the *explicit FLP check*, constructs $f\Pi^{\mathbf{A}}$ and checks whether it has some model smaller than \mathbf{A} . However, this approach is not very efficient in practice; actually the explicit FLP check often dominates total runtime. This calls for an alternative method to perform an FLP check efficiently, which we address in this paper.

For ordinary logic programs, *unfounded sets* proved to be a fruitful approach for minimality checking [Leone et al., 1997]. Informally, a set U of atoms is unfounded wrt. an interpretation \mathbf{A} , if no atom in U is derivable by applying a rule if all atoms in U are made false; this in particular allows to identify mutual and self-support by rules. E.g. $U = \{a, b\}$ is unfounded with respect to $\mathbf{A} = \{\mathbf{T}a, \mathbf{T}b\}$ (a and b are true) under the rules $a \leftarrow b$ and $b \leftarrow a$. The answer sets of a program are now the interpretations that are unfounded-free, i.e., disjoint from all their unfounded sets. Thus to check a candidate answer set, one search for unfounded sets of Π wrt. \mathbf{A} rather than explicitly construct $f\Pi^{\mathbf{A}}$ and enumerate its models in search for a smaller one.

Starting from this idea, we define unfounded sets for HEX-programs following [Faber, 2005], which extended the concept to programs with aggregates, and explore their use for FLP checking. Our main contributions are summarized as follows:

- We present a basic encoding of unfounded set existence to a set of *nogoods*, i.e., constraints that have to be satisfied, and we show that its solutions are in 1-1 correspondence with the unfounded sets. The latter

can thus be computed running a SAT solver and followed by a post-processing step which checks that the values of replacement atoms are compatible with the external call results. Benchmarks show that this strategy is already more efficient than the explicit FLP check.

- We then present an advanced encoding of unfounded set existence that is independent of the interpretation. Compared to our first encoding, it is conceptually more involved and has (slightly) higher initialization cost, but it has the advantage that it can be reused for all unfounded set checks and needs no separate initialization for each check. Our benchmarks show that the advanced encoding is superior to the first one for many practical problems.
- Next, we consider optimizations which hinge on dependencies between external and ordinary atoms that are determined in careful analysis. These optimizations can be integrated into our encodings by adding further nogoods which restrict the search space to relevant parts.
- We consider how to exploit information gained in the FLP check for generating candidate answer sets in step 1. Adopting a Conflict Driven Clause Learning approach [Drescher et al., 2008], this step has been recently enhanced by *external behavior learning* [Eiter et al., 2012a], in which nogoods describing the external source behavior are learned during the search to guide model generation towards proper guesses. We show how to learn in step 1 additional nogoods from unfounded sets that avoid the reconstruction of the same or related unfounded sets, yielding further gain.
- We present a *syntactic decision criterion* that can be used to decide whether a program possibly has unfounded sets. If the result of this check is negative, then the computationally expensive search for unfounded sets can be skipped entirely. The criterion is based on atom dependency and, loosely speaking, says that there are no cyclic dependencies of ground atoms through external atoms. This property can be efficiently checked for a given ground HEX-program using standard methods. In fact it applies to a range of applications, in particular for input-stratified programs, where external sources are accessed in a workflow to produce input for the next stage of computation. However, advanced applications of HEX-programs can have cycles through external atoms, e.g., in natural encodings of problems on multi-context systems [Brewka and Eiter, 2007] or abstract argumentation systems [Dung, 1995], for which the FLP check can not be simply skipped.
- In further elaboration, we then consider a *program decomposition* based on the dependency graph that is induced by a program Π . We show that Π has some unfounded set wrt. a candidate answer set \mathbf{A} exactly if some of the components Π_C in its decomposition has some unfounded set wrt. \mathbf{A} ; As computing the decomposition is efficiently possible and does not incur large overhead, we can apply the decision criterion for skipping the FLP check efficiently on a finer-grained level, and the search for unfounded sets can be guided to relevant program parts.
- An experimental evaluation on advanced reasoning applications shows that unfounded sets checking combined with learning methods from [Eiter et al., 2012a] improves HEX-program evaluation considerably, sometimes drastically. More specifically, the benchmark applications include reasoning tasks in Multi-Context Systems [Brewka and Eiter, 2007, Eiter et al., 2012e], abstract argumentation [Dung, 1995], terminological default reasoning over description logic knowledge bases [Baader and Hollunder, 1995], and conformant planning [Goldman and Boddy, 1996, Smith and Weld, 1998]. The decomposition approach can yield a considerable gain, as it appears e.g. for the HEX-encoding of a Dung-style 1995 argumentation semantics and DL-programs [Eiter et al., 2008a]. On the other hand, for the terminological default reasoning benchmark, our syntactic criterion let us conclude that the FLP check is obsolete. In conclusion, the new approach enables significant speedup and thus enlarges the scope of HEX applicability.

Organization. The rest of this paper is organized as follows. The next section provides preliminaries on HEX programs and their evaluation via compatible sets. In Section 3, we define unfounded sets and present a basic and a uniform encoding of unfounded set search using nogoods. Section 4 considers refinements and optimizations of the encodings, as well as external behavior learning to prevent reconstruction of unfounded sets. In Section 5, we give a syntactic decision criterion to avoid the FLP check and a program decomposition method exploiting it. Experimental results of a prototype implementation are reported in Section 6. In Section 7 we consider related work and extensions of our approach. In Section 8 we conclude and point out issues for further research.

2 Preliminaries

In this section, we start with some basic definitions, and then introduce HEX-programs.

In accordance with [Gebser et al., 2012, Eiter et al., 2012a], a (*signed*) *literal* is a positive or a negative formula $\mathbf{T}a$ resp. $\mathbf{F}a$, where a is a ground atom of form $p(c_1, \dots, c_\ell)$, with predicate p and constants c_1, \dots, c_ℓ , abbreviated $p(\mathbf{c})$. For a literal $\sigma = \mathbf{T}a$ or $\sigma = \mathbf{F}a$, let $\bar{\sigma}$ denote its opposite, i.e., $\overline{\mathbf{T}a} = \mathbf{F}a$ and $\overline{\mathbf{F}a} = \mathbf{T}a$. An *assignment* \mathbf{A} over a (finite) set of atoms A is a consistent set of signed literals $\mathbf{T}a$ or $\mathbf{F}a$, where $\mathbf{T}a$ expresses that $a \in A$ is true and $\mathbf{F}a$ that it is false; \mathbf{A} is *complete*, also called an *interpretation*, if no assignment $\mathbf{A}' \supset \mathbf{A}$ exists. We denote by $\mathbf{A}^{\mathbf{T}} = \{a \mid \mathbf{T}a \in \mathbf{A}\}$ and $\mathbf{A}^{\mathbf{F}} = \{a \mid \mathbf{F}a \in \mathbf{A}\}$ the set of atoms that are true resp. false in \mathbf{A} , and by $\text{ext}(q, \mathbf{A}) = \{\mathbf{c} \mid \mathbf{T}q(\mathbf{c}) \in \mathbf{A}\}$ the extension of a predicate q in \mathbf{A} . Furthermore, $\mathbf{A}|_q$ is the set of all literals over atoms with predicate q in \mathbf{A} . For a list $\mathbf{q} = q_1, \dots, q_k$ of predicates, we write $p \in \mathbf{q}$ iff $q_i = p$ for some $1 \leq i \leq k$, and let $\mathbf{A}|_{\mathbf{q}} = \bigcup_j \mathbf{A}|_{q_j}$.

A *nogood* is a set $\{L_1, \dots, L_n\}$ of literals $L_i, 1 \leq i \leq n$. An interpretation \mathbf{A} is a *solution* to a nogood δ (resp. to a set Δ of nogoods), iff $\delta \not\subseteq \mathbf{A}$ (resp. $\delta \not\subseteq \mathbf{A}$ for all $\delta \in \Delta$).

Example 2.1. *The interpretation $\mathbf{A} = \{\mathbf{T}a, \mathbf{F}b\}$ is a solution to the nogood $\{\mathbf{T}a, \mathbf{T}b, \mathbf{T}c\}$ but not to $\{\mathbf{T}a, \mathbf{F}b, \mathbf{T}c\}$.*

2.1 HEX Programs

HEX-Program Syntax. As introduced in [Eiter et al., 2005], HEX-programs are a generalization of (disjunctive) extended logic programs under the answer set semantics [Gelfond and Lifschitz, 1991]; for details and background see [Eiter et al., 2005]. HEX-programs extend ordinary ASP programs by *external atoms*, which enable a bidirectional interaction between a program and external sources of computation. External atoms have a list of input parameters (constants or predicate names) and a list of output parameters. Informally, to evaluate an external atom, the reasoner passes the constants and extensions of the predicates in the input tuple to the external source associated with the external atom. The external source computes output tuples which are matched with the output list. Syntactically, a *ground external atom* is of the form

$$\&g[\mathbf{p}](\mathbf{c}), \tag{1}$$

where $\mathbf{p} = p_1, \dots, p_k$ are constant input parameters (predicate names or object constants), and $\mathbf{c} = c_1, \dots, c_\ell$ are constant output terms.

A *default literal* is a formula b or *not* b , where b is a (ground) ordinary or external atom. For every set S of ordinary and external atoms, we let $\text{not } S = \{\text{not } b \mid b \in S\}$.

Ground HEX-programs are then defined similar to ground ordinary ASP programs.

Definition 2.2 (Ground HEX-programs). A ground HEX-program consists of rules

$$a_1 \vee \dots \vee a_k \leftarrow b_1, \dots, b_m, \text{not } b_{m+1}, \dots, \text{not } b_n, \quad (2)$$

where each a_i is an (ordinary) ground atom $p(c_1, \dots, c_\ell)$ with constants c_i , $1 \leq i \leq \ell$, each b_j is either an ordinary ground atom or a ground external atom, and $k + n > 0$.¹

For a rule r , the *head* is given as $H(r) = \{a_1, \dots, a_k\}$ and the *body* is $B(r) = B^+(r) \cup \text{not } B^-(r)$, where $B^+(r) = \{b_1, \dots, b_m\}$ is the *positive body*, $B^-(r) = \{b_{m+1}, \dots, b_n\}$ is the *negative body*. If $B(r) = \emptyset$, then r is a *fact*, and we omit \leftarrow . For a program Π , let $A(\Pi)$ be the set of all ordinary atoms occurring in Π . For a default literal b , let $\mathbf{t}b = \mathbf{T}a$ if $b = a$ for an atom a , and $\mathbf{t}b = \mathbf{F}a$ if $b = \text{not } a$. Furthermore, $\mathbf{f}b = \mathbf{F}a$ if $b = a$ and $\mathbf{f}b = \mathbf{T}a$ if $b = \text{not } a$.

Example 2.3. For rule $r = a \vee b \leftarrow c, \text{not } d$ we have $H(r) = \{a, b\}$, $B^+(r) = \{c\}$ and $B^-(r) = \{d\}$. We further have $\mathbf{t}c = \mathbf{T}c$, $\mathbf{f}c = \mathbf{F}c$, $\mathbf{t} \text{not } d = \mathbf{F}d$ and $\mathbf{f} \text{not } d = \mathbf{T}d$.

We also use non-ground programs (i.e., with variables allowed in place of object constants). For such programs, suitable safety conditions allow for using a grounding procedure [Eiter et al., 2006]. We limit our investigation to ground programs.

Example 2.4. The program

$$\begin{aligned} & \text{domain}(a); \quad \text{domain}(b) \\ & \quad \text{sel}(X) \leftarrow \text{domain}(X), \text{not } \text{n sel}(X) \\ & \quad \text{n sel}(X) \leftarrow \text{domain}(X), \text{not } \text{sel}(X) \end{aligned}$$

partitions the domain elements into two sets *sel* and *n sel*.

HEX-Program Semantics. An ordinary ground atom a is *true* relative to assignment \mathbf{A} , denoted $\mathbf{A} \models a$, if $\mathbf{T}a \in \mathbf{A}$ and *false* otherwise. A default-negated ground atom $\text{not } a$ is *true* relative to assignment \mathbf{A} , denoted $\mathbf{A} \models \text{not } a$, if $\mathbf{F}a \in \mathbf{A}$ and *false* otherwise.

The semantics of a ground external atom $\&g[\mathbf{p}](\mathbf{c})$ wrt. an interpretation \mathbf{A} is given by the value of a $1+k+l$ -ary Boolean *oracle function* $f_{\&g}$ that is defined for all possible values of \mathbf{A} , \mathbf{p} and \mathbf{c} , such that $\&g[\mathbf{p}](\mathbf{c})$ is true relative to \mathbf{A} if and only if $f_{\&g}(\mathbf{A}, \mathbf{p}, \mathbf{c}) = 1$.

Example 2.5 (Set Partitioning). Consider the program Π

$$\begin{aligned} & \text{sel}(a) \leftarrow \text{domain}(a), \&diff[\text{domain}, \text{n sel}](a) \\ & \text{n sel}(a) \leftarrow \text{domain}(a), \&diff[\text{domain}, \text{sel}](a) \\ & \text{domain}(a) \end{aligned}$$

where $\&diff[p, q](X)$ computes the set of all elements X which are in the extension of p but not in the extension of q . Informally, this program implements a choice from $\text{sel}(a)$ and $\text{n sel}(a)$.

¹ For simplicity, we do not formally introduce strong negation but view, as customary, classical literals $\neg a$ as new atoms together with a constraint which disallows that a and $\neg a$ are simultaneously true.

Satisfaction of ordinary rules and ASP programs [Gelfond and Lifschitz, 1991] is then extended to HEX-rules and programs in the obvious way: a rule r is satisfied by assignment \mathbf{A} , denoted $\mathbf{A} \models r$, iff $\mathbf{A} \models h$ for some $h \in H(r)$, or $\mathbf{A} \not\models b$ for some $b \in B^+(r)$, or $\mathbf{A} \models b$ for some $b \in B^-(r)$. A program Π is satisfied by assignment \mathbf{A} iff $\mathbf{A} \models r$ for all $r \in \Pi$.

The notion of extension $ext(\cdot, \mathbf{A})$ for external predicates $\&g$ with input lists \mathbf{p} is naturally defined by $ext(\&g[\mathbf{p}], \mathbf{A}) = \{\mathbf{c} \mid f_{\&g}(\mathbf{A}, \mathbf{p}, \mathbf{c}) = 1\}$.

An input predicate p of an external predicate with input list $\&g[\mathbf{p}]$ is *monotonic* (*antimonotonic*), iff $f_{\&g}(\mathbf{A}, \mathbf{p}, \mathbf{c}) = 1$ implies $f_{\&g}(\mathbf{A}', \mathbf{p}, \mathbf{c}) = 1$ ($f_{\&g}(\mathbf{A}, \mathbf{p}, \mathbf{c}) = 0$ implies $f_{\&g}(\mathbf{A}', \mathbf{p}, \mathbf{c}) = 0$) for all \mathbf{A}' s.t. $ext(p, \mathbf{A}') \supseteq ext(p, \mathbf{A})$ and $ext(q, \mathbf{A}') = ext(q, \mathbf{A})$ for $q \in \mathbf{p}$ and $q \neq p$. The sublist of all monotonic resp. antimonotonic p is denoted by $m(\&g)$ resp. $a(\&g)$ and the sublist of other (i.e., *nonmonotonic*) p by $n(\&g)$; we also write \mathbf{p}_τ for $\tau(\&g)$, $\tau \in \{m, a, n\}$.

Definition 2.6 (FLP-Reduct [Faber et al., 2011]). *For an interpretation \mathbf{A} over a program Π , the FLP-reduct of Π wrt. \mathbf{A} is the set $f\Pi^{\mathbf{A}} = \{r \in \Pi \mid \mathbf{A} \models b, \text{ for all } b \in B(r)\}$ of all rules whose body is satisfied under \mathbf{A} .*

An assignment \mathbf{A}_1 is smaller or equal to another assignment \mathbf{A}_2 wrt. a program Π , denoted $\mathbf{A}_1 \leq_\Pi \mathbf{A}_2$, iff $\{\mathbf{T}a \in \mathbf{A}_1 \mid a \in A(\Pi)\} \subseteq \{\mathbf{T}a \in \mathbf{A}_2 \mid a \in A(\Pi)\}$.

Definition 2.7 (Answer Set). *An answer set of Π is a \leq_Π -minimal model \mathbf{A} of $f\Pi^{\mathbf{A}}$.*

Since interpretations (thus answer sets) are complete assignments, in slight abuse of notation we identify them with their sets of positive literals.

Example 2.8. *Consider the program Π :*

$$\begin{aligned} p &\leftarrow \&id[q]() \\ q &\leftarrow p \end{aligned}$$

where $f_{\&id}(\mathbf{A}, p, \cdot) = 1$ iff $\mathbf{T}p \in \mathbf{A}$ is true. Then Π has the answer set $\mathbf{A}_1 = \emptyset$; indeed it is a \leq_Π -minimal model of $f\Pi^{\mathbf{A}_1} = \emptyset$. Note that $\mathbf{A}_2 = \{p, q\}$ is not an answer set of Π , as it is not a minimal model of $f\Pi^{\mathbf{A}_2} = \Pi$

2.2 Evaluation

The answer sets of a HEX-program Π are determined by the dlhex solver using a transformation to ordinary ASP programs as follows. Each external atom $\&g[\mathbf{p}](\mathbf{c})$ in Π is replaced by an ordinary ground *replacement atom* $e_{\&g[\mathbf{p}]}(\mathbf{c})$ and a rule $e_{\&g[\mathbf{p}]}(\mathbf{c}) \vee ne_{\&g[\mathbf{p}]}(\mathbf{c}) \leftarrow$ is added to the program. The answer sets of the resulting *guessing program* $\hat{\Pi}$ are determined by an ordinary ASP solver and projected to non-replacement atoms. However, the resulting assignments are not necessarily models of Π , as the value of $\&g[\mathbf{p}]$ under $f_{\&g}$ can be different from the one of $e_{\&g[\mathbf{p}]}(\mathbf{c})$. Each answer set of $\hat{\Pi}$ is thus merely a *candidate* which must be checked against the external sources. If no discrepancy is found, the model candidate is a *compatible set* of Π . More precisely,

Definition 2.9 (Compatible Set). *A compatible set of a program Π is an assignment $\hat{\mathbf{A}}$ such that*

- (i) $\hat{\mathbf{A}}$ is an answer set [Gelfond and Lifschitz, 1991] of the guessing program $\hat{\Pi}$, and
- (ii) $f_{\&g}(\hat{\mathbf{A}}, \mathbf{p}, \mathbf{c}) = 1$ iff $\mathbf{T}e_{\&g[\mathbf{p}]}(\mathbf{c}) \in \hat{\mathbf{A}}$ for all external atoms $\&g[\mathbf{p}](\mathbf{c})$ in Π , i.e. the guessed values coincide with the actual output under the input from $\hat{\mathbf{A}}$.

The compatible sets of Π computed by `dlvhex` include (modulo $A(\Pi)$) all answer sets of Π , where each answer set \mathbf{A} of Π is the restriction of a unique compatible set $\hat{\mathbf{A}}$ to non-replacement atoms. To filter out the compatible sets not yielding answer sets, the current evaluation algorithm proceeds as follows. Each compatible set $\hat{\mathbf{A}}$ is fed to the *FLP check*, which constructs the reduct $f\Pi^{\hat{\mathbf{A}}}$ and checks whether it has a model \mathbf{A}' smaller than \mathbf{A} , by enumerating the models \mathbf{A}' of $f\Pi^{\hat{\mathbf{A}}}$ and comparing them to \mathbf{A} ; if such an \mathbf{A}' is found, it rejects \mathbf{A} , otherwise outputs \mathbf{A} as an answer set.

The explicit FLP check also uses a rewriting to ordinary programs and corresponds to the search for compatible sets of the following program:

$$\begin{aligned} \text{Check}(\Pi, \mathbf{A}) = & f\hat{\Pi}^{\hat{\mathbf{A}}} \cup \{\leftarrow a \mid a \in A(\Pi), \mathbf{T}a \notin \hat{\mathbf{A}}\} \cup \{a \vee a' \leftarrow \mid \mathbf{T}a \in \hat{\mathbf{A}}\} \\ & \cup \{\leftarrow \text{not smaller}\} \cup \{\text{smaller} \leftarrow \text{not } a \mid a \in A(\Pi), \mathbf{T}a \in \hat{\mathbf{A}}\}. \end{aligned}$$

It consists of the reduct $f\hat{\Pi}^{\hat{\mathbf{A}}}$ and rules that restrict the search to proper subinterpretations of $\hat{\mathbf{A}}$, where *smaller* is a new atom. Moreover, as we actually need to search for models and not just compatible sets, rules of the form $a \vee a' \leftarrow$ (where a' is a new atom for each $\mathbf{T}a \in \hat{\mathbf{A}}$) make sure that atoms can be arbitrarily true without having a justifying rule in Π . Because of these guessing rules, we can rewrite all rules in $f\hat{\Pi}^{\hat{\mathbf{A}}}$ except the guesses on replacement atoms to constraints; this is more efficient for evaluation. Our comparison in Section 6 uses this optimized version of the explicit check, but still demonstrates a significant performance gain by our novel approach.

Example 2.10 (cont'd). *Reconsider the program $\Pi = \{p \leftarrow \&id[q](); q \leftarrow p\}$ from above. Then the corresponding guessing program is $\hat{\Pi} = \{p \leftarrow e_{\&id[q]}(); q \leftarrow p; e_{\&id[q]} \vee ne_{\&id[q]} \leftarrow\}$ and yields the compatible sets $\hat{\mathbf{A}}_1 = \emptyset$ and $\hat{\mathbf{A}}_2 = \{\mathbf{T}p, \mathbf{T}q, \mathbf{T}e_{\&id[p]}\}$. While $\mathbf{A}_1 = \emptyset$ is also a \leq_{Π} -minimal model of $f\Pi^{\hat{\mathbf{A}}_1} = \emptyset$, $\mathbf{A}_2 = \{\mathbf{T}p, \mathbf{T}q\}$ is not a \leq_{Π} -minimal model of $f\Pi^{\hat{\mathbf{A}}_2} = \Pi$.*

3 Unfounded Set Detection

As described in Section 2.2, the explicit minimality check is computationally costly and involves much overhead: all models of $\text{Check}(\Pi, \mathbf{A})$ must be enumerated, and calls to the external sources to test compatibility must be made. Even worse, as we need to search for a smaller *model* and not just for a smaller compatible set, $\text{Check}(\Pi, \mathbf{A})$ usually (to our experience) has even more models than the original program. Moreover, it appears that in many current application scenarios there is no smaller model of the reduct $f\Pi^{\hat{\mathbf{A}}}$, i.e., most assignments \mathbf{A} extracted from compatible sets $\hat{\mathbf{A}}$ pass the FLP check. There are two possible reliefs: developing a cheaper minimality check and to avoid the minimality check if possible. While this section targets at the former idea, the latter one is addressed in Section 5.

To this end, we present a novel FLP check algorithm based on unfounded sets (UFS). Instead of explicitly searching for smaller models of the reduct, we check whether the candidate answer set is unfounded-free. To this end, we use unfounded sets for hex programs akin to those in [Faber, 2005] for programs with arbitrary aggregates.

Definition 3.1 (Unfounded Set). *Given a program Π and an assignment \mathbf{A} , let X be any set of ordinary ground atoms appearing in Π . Then, X is an unfounded set for Π wrt. \mathbf{A} if, for each rule r having some atoms from X in the head, at least one of the following conditions holds, where $\mathbf{A} \dot{\cup} \neg.X = (\mathbf{A} \setminus \{\mathbf{T}a \mid a \in X\}) \cup \{\mathbf{F}a \mid a \in X\}$:*

- (i) *some literal of $B(r)$ is false wrt. \mathbf{A} ,*

- (ii) some literal of $B(r)$ is false wrt. $\mathbf{A} \dot{\cup} \neg.X$, or
- (iii) some atom of $H(r) \setminus X$ is true wrt. \mathbf{A} .

Intuitively, an unfounded set is a set of atoms which only cyclically support each other; by assigning all of them false, no violation of any rule will be introduced. As for answer sets, their minimality enforces now that no subset of the atoms that are true in an answer set can form an unfounded set; in fact, answer sets can be characterized in terms of unfounded sets, using the following notion.

Definition 3.2 (Unfounded-free Interpretations). *An interpretation \mathbf{A} of a program Π is unfounded-free, iff $\mathbf{A}^T \cap X = \emptyset$, for every unfounded set X of Π wrt. \mathbf{A} .*

The following result is a generalization of a respective result for ordinary (disjunctive) logic programs [Leone et al., 1997] and logic programs with aggregates [Faber et al., 2011].

Theorem 3.3 (Characterization of Answer Sets). *A model \mathbf{A} of a program Π is an answer set of Π iff it is unfounded-free.*

Example 3.4. *Consider the program Π and \mathbf{A}_1 from Example 2.10. Trivially, \mathbf{A}_1 is unfounded-free, and thus \mathbf{A}_1 is an answer set of Π . On the other hand, the set $X = \{p, q\}$ is an unfounded set w.r.t. \mathbf{A}_2 , since X intersects with the head of $p \leftarrow \&id[q]()$ and $\mathbf{A} \dot{\cup} \neg.X \not\models \&id[q]()$. Therefore \mathbf{A}_2 is not unfounded-free and not an answer set.*

3.1 Basic Encoding of the Unfounded Set Search

We realize the search for unfounded sets using nogoods, i.e., for a given Π and an assignment \mathbf{A} we construct a set of nogoods, such that solutions to this set correspond to (potential) unfounded sets; we then use a SAT solver to search for such unfounded sets.

More specifically, our encoding of unfounded set detection uses a set

$$\Gamma_{\Pi}^{\mathbf{A}} = N_{\Gamma, \Pi}^{\mathbf{A}} \cup O_{\Gamma, \Pi}^{\mathbf{A}},$$

of nogoods where $N_{\Gamma, \Pi}^{\mathbf{A}}$ contains all *necessary* constraints and the set $O_{\Gamma, \Pi}^{\mathbf{A}}$ are *optional* optimization nogoods that prune irrelevant parts of the search space; the latter is related to [Drescher et al., 2008] but respects external atoms. The idea is that the set of ordinary atoms of a solution to $\Gamma_{\Pi}^{\mathbf{A}}$ represents a (potential) unfounded set U of Π wrt. \mathbf{A} , while the replacement atoms encode the truth values of the corresponding external atoms under $\mathbf{A} \dot{\cup} \neg.U$.

Let $B_o^+(r) \subseteq B^+(r)$ consist of all ordinary atoms, and let $B_e(r) \subseteq B(r)$ consist of all external replacement atoms. Then, $\Gamma_{\Pi}^{\mathbf{A}}$ is built over atoms $A(\Gamma_{\Pi}^{\mathbf{A}}) = A(\hat{\Pi}) \cup \{h_r, l_r \mid r \in \Pi\}$, where h_r , and l_r are atoms for every rule r in Π . The *necessary part*

$$N_{\Gamma, \Pi}^{\mathbf{A}} = \{\{\mathbf{F}a \mid \mathbf{T}a \in \mathbf{A}\}\} \cup \left(\bigcup_{r \in \Pi} R_{r, \mathbf{A}} \right)$$

of $\Gamma_{\Pi}^{\mathbf{A}}$ consists of a nogood $\{\mathbf{F}a \mid \mathbf{T}a \in \mathbf{A}\}$, which eliminates unfounded sets that do not intersect with true atoms in \mathbf{A} , and of nogoods $R_{r, \mathbf{A}} = H_{r, \mathbf{A}} \cup C_{r, \mathbf{A}}$ for every $r \in \Pi$ where

- $H_{r, \mathbf{A}} = \{\{\mathbf{T}h_r\} \cup \{\mathbf{F}h \mid h \in H(r)\}\} \cup \{\{\mathbf{F}h_r, \mathbf{T}h\} \mid h \in H(r)\}$, called the *head criterion*, encodes that h_r is true for a rule r iff some atom of $H(r)$ is in the unfounded set; and

$$\bullet C_{r,\mathbf{A}} = \begin{cases} \{\{\mathbf{T}h_r\} \cup \\ \{\mathbf{F}a \mid a \in B_o^+(r), \mathbf{A} \models a\} \cup \{\mathbf{t}a \mid a \in B_e(\hat{r})\} \cup \\ \{\mathbf{T}h \mid h \in H(r), \mathbf{A} \models h\}\} & \text{if } \mathbf{A} \models B(r), \\ \{\} & \text{otherwise,} \end{cases}$$

called the *conditional part* $C_{r,\mathbf{A}}$, encodes that Condition (i), (ii) or (iii) of Definition 3.1 must hold if h_r is true.

More specifically, for an unfounded set U and a rule r with $H(r) \cap U \neq \emptyset$ (h_r is true) it must not happen that $\mathbf{A} \models B(r)$ (Condition (i) fails), no $a \in B_o^+(r)$ with $\mathbf{A} \models a$ is in the unfounded set and all $a \in B_e(\hat{r})$ are true under $\mathbf{A} \cup \neg.U$ (Condition (ii) fails), and all $h \in H(r)$ with $\mathbf{A} \models h$ are in the unfounded set (Condition (iii) fails). Concrete instances for $O_{\Pi}^{\Gamma,\mathbf{A}}$ are defined in Section 4.

Example 3.5. Consider $\Pi = \{r_1 : p \leftarrow \&id[p]()\}$ and the compatible set $\hat{\mathbf{A}} = \{\mathbf{T}p, \mathbf{T}e_{\&id[p]}\}$. The nogood set $N_{\Gamma,\Pi}^{\mathbf{A}_2}$ is $\{\{\mathbf{T}h_{r_1}, \mathbf{F}p\}, \{\mathbf{F}h_{r_1}, \mathbf{T}p\}, \{\mathbf{T}h_{r_1}, \mathbf{T}e_{\&id[p]}\}, \{\mathbf{T}p\}\}$.

Towards computing unfounded sets, observe that every unfounded set can be extended to a solution to the nogood set $\Gamma_{\Pi}^{\mathbf{A}}$ over $A(\Gamma_{\Pi}^{\mathbf{A}})$. Conversely, the solutions to $\Gamma_{\Pi}^{\mathbf{A}}$ include specific extensions of the unfounded sets, given for each unfounded set U by assigning true to all atoms in U , to all h_r such that $H(r)$ intersects with U , and to all replacement atoms $e_{\&g[p]}(\mathbf{c})$ such that $\&g[p](\mathbf{c})$ is true under $\mathbf{A} \cup \neg.U$, and assigning false to all other atoms in $A(\Gamma_{\Pi}^{\mathbf{A}})$. More formally,

Definition 3.6 (Induced Assignment of an Unfounded Set wrt. $\Gamma_{\Pi}^{\mathbf{A}}$). Let U be an unfounded set of a program Π wrt. assignment \mathbf{A} . The assignment induced by U wrt. $\Gamma_{\Pi}^{\mathbf{A}}$, denoted $I_{\Gamma}(U, \Gamma_{\Pi}^{\mathbf{A}})$, is

$$I_{\Gamma}(U, \Gamma_{\Pi}^{\mathbf{A}}, \Pi, \mathbf{A}) = I_{\Gamma}^0(U, \Pi, \mathbf{A}) \cup \{\mathbf{F}a \mid a \in A(\Gamma_{\Pi}^{\mathbf{A}}), \mathbf{T}a \notin I_{\Gamma}^0(U, \Pi, \mathbf{A})\},$$

where

$$I_{\Gamma}^0(U, \Pi, \mathbf{A}) = \{\mathbf{T}a \mid a \in U\} \cup \{\mathbf{T}h_r \mid r \in \Pi, H(r) \cap U \neq \emptyset\} \cup \{\mathbf{T}e_{\&g[p]}(\mathbf{c}) \mid e_{\&g[p]}(\mathbf{c}) \in A(\hat{\Pi}), \mathbf{A} \cup \neg.U \models \&g[p](\mathbf{c})\}.$$

Intuitively, the truth value of a in the assignment represents whether a is in the unfounded set U or not, the h_r encode if the intersection of $H(r)$ with U is nonempty, and the replacement atoms represent the truth values of the corresponding external atoms under $\mathbf{A} \cup \neg.U$.

We call a set N of nogoods *conservative*, if it holds for every unfounded set U of Π wrt. \mathbf{A} that $I_{\Gamma}(U, \Gamma_{\Pi}^{\mathbf{A}}, \Pi, \mathbf{A})$ is a solution to N . We then show that the solutions to $\Gamma_{\Pi}^{\mathbf{A}}$ include all assignments induced by unfounded sets of Π wrt. \mathbf{A} , assuming that $O_{\Pi}^{\Gamma,\mathbf{A}}$ is conservative.

Proposition 3.7. Let U be an unfounded set of a program Π wrt. assignment \mathbf{A} such that $\mathbf{A}^{\mathbf{T}} \cap U \neq \emptyset$. Then $I_{\Gamma}(U, \Gamma_{\Pi}^{\mathbf{A}}, \Pi, \mathbf{A})$ is a solution to $\Gamma_{\Pi}^{\mathbf{A}}$.

Note that the converse does not hold, i.e., not every solution corresponds to some induced assignment; intuitively this is because it does not reflect the semantics of external sources. Regardless of this we immediately obtain from Proposition 3.7 a useful test for unfounded-freeness.

Corollary 3.8. If $\Gamma_{\Pi}^{\mathbf{A}}$ has no solution, then $U \cap \mathbf{A}^{\mathbf{T}} = \emptyset$ for every unfounded set U of Π .

Using the following result, we can find the unfounded sets of Π wrt. \mathbf{A} among all solutions to $\Gamma_{\Pi}^{\mathbf{A}}$ by using a postcheck on the external atoms.

Theorem 3.9. Let S be a solution to $\Gamma_{\Pi}^{\mathbf{A}}$ such that

- (a) $\mathbf{T}e_{\&g[\mathbf{p}]}(\mathbf{c}) \in S$ and $\mathbf{A} \not\models \&g[\mathbf{p}](\mathbf{c})$ implies $\mathbf{A} \dot{\cup} \neg.U \models \&g[\mathbf{p}](\mathbf{c})$; and
(b) $\mathbf{F}e_{\&g[\mathbf{p}]}(\mathbf{c}) \in S$ and $\mathbf{A} \models \&g[\mathbf{p}](\mathbf{c})$ implies $\mathbf{A} \dot{\cup} \neg.U \not\models \&g[\mathbf{p}](\mathbf{c})$

where $U = \{a \mid a \in A(\Pi), \mathbf{T}a \in S\}$. Then U is an unfounded set of Π wrt. \mathbf{A} .

Informally, the proposition states that the non-replacement atoms in S that are true and also appear in Π form an unfounded set, provided that truth of the replacement atoms $e_{\&g[\mathbf{p}]}(\mathbf{c})$ in S coincides with the truth of the corresponding $\&g[\mathbf{p}](\mathbf{c})$ under $\mathbf{A} \dot{\cup} \neg.U$ (as in Definition 3.6). However, this check is just required if the truth values of $e_{\&g[\mathbf{p}]}(\mathbf{c})$ in S and of $\&g[\mathbf{p}](\mathbf{c})$ under \mathbf{A} differ. This gives rise to an important optimization for the implementation: external atoms, whose (known) truth value of $\&g[\mathbf{p}](\mathbf{c})$ under \mathbf{A} matches the truth value of $e_{\&g[\mathbf{p}]}(\mathbf{c})$ in S , do not need to be postchecked.

Example 3.10. Reconsider program $\Pi = \{r_1 : p \leftarrow \&id[p]()\}$ from Example 3.5 and the compatible set $\hat{\mathbf{A}}_2 = \{\mathbf{T}p, \mathbf{T}e_{\&id[p]}\}$. The nogood set $N_{\Gamma, \Pi}^{\hat{\mathbf{A}}_2} = \{\{\mathbf{T}h_{r_1}, \mathbf{F}p\}, \{\mathbf{F}h_{r_1}, \mathbf{T}p\}, \{\mathbf{T}h_{r_1}, \mathbf{T}e_{\&id[p]}(), \mathbf{T}p\}\}$ has solutions $S \supseteq \{\mathbf{T}h_{r_1}, \mathbf{T}p, \mathbf{F}e_{\&id[p]}()\}$, which correspond to the unfounded set $U = \{p\}$. Here, $\mathbf{F}e_{\&id[p]}()$ represents that $\mathbf{A}_2 \dot{\cup} \neg.U \not\models \&id[p]()$.

Note that due to the premises in Conditions (a) and (b) of Theorem 3.9, the postcheck is faster if $\mathbf{T}e_{\&g[\mathbf{p}]}(\mathbf{c}) \in S$ whenever $\mathbf{A} \models \&g[\mathbf{p}](\mathbf{c})$ holds for many external atoms in Π . This can be exploited during the construction of S as follows: if it is not absolutely necessary to set the truth value of $e_{\&g[\mathbf{p}]}(\mathbf{c})$ differently, then carry over the value from $\&g[\mathbf{p}](\mathbf{c})$ under \mathbf{A} . Specifically, this is successful if $e_{\&g[\mathbf{p}]}(\mathbf{c})$ does not occur in $\Gamma_{\Pi}^{\hat{\mathbf{A}}}$.

3.2 Uniform Encoding of the Unfounded Set Search

The encoding $\Gamma_{\Pi}^{\hat{\mathbf{A}}}$ presented in the previous subsection has the disadvantage that it depends on the current assignment \mathbf{A} . Therefore it needs to be generated separately for every unfounded set check if the assignment changed (which is very likely). As this causes significant overhead, we present now an advanced encoding which is independent of the current assignment. We introduce some additional variables which represent the truth values of the atoms in the current assignment. Prior to an unfounded set check, the current assignment is injected by setting the values of these variables to fixed values, which can be done using *assumptions* as supported by modern SAT solvers such as CLASP. Changing assumptions is much easier than changing the encoding, which leads to an additional speedup in some cases, especially for programs which need many unfounded set checks.

Our advanced encoding uses a set Ω_{Π} of nogoods. As before, the idea is that the set of ordinary atoms of a solution to Ω_{Π} represents a (potential) unfounded set U of Π wrt. some assignment \mathbf{A} , while the replacement atoms encode the truth values of the corresponding external atoms under $\mathbf{A} \dot{\cup} \neg.U$. The encoding Ω_{Π} is conceptually more complex than $\Gamma_{\Pi}^{\hat{\mathbf{A}}}$; the initialization is computationally (slightly) more costly, hence the advantages of our new encoding become visible for instances with many compatible sets (thus many unfounded set checks), while it might be counterproductive for small instances.

The nogood set Ω_{Π} is built over the atoms of $\hat{\Pi}$ and further fresh atoms not occurring in $\hat{\Pi}$: h_r and l_r , for every rule r in Π , $a_{\mathbf{A}}$ for every ordinary atom $a \in A(\hat{\Pi})$ (i.e. ordinary atoms in Π and replacement atom auxiliaries), and $a_{\mathbf{A} \dot{\cup} \neg.U}$, $a_{\mathbf{A} \wedge U}$, $a_{\overline{\mathbf{A}} \vee U}$ for every ordinary atom $a \in A(\Pi)$. The *auxiliary atoms* $a_{\mathbf{A}}$, $a_{\mathbf{A} \dot{\cup} \neg.U}$, $a_{\mathbf{A} \wedge U}$, $a_{\overline{\mathbf{A}} \vee U}$ are used to make the encoding independent of a certain assignment \mathbf{A} . Only during the unfounded set check with respect to a certain assignment, we will temporarily add assumptions to the solver which force certain truth values of the atoms $a_{\mathbf{A}}$ for all $a \in A(\hat{\Pi})$ depending on the current assignment \mathbf{A} .

Intuitively, $a_{\mathbf{A}}$ represents the truth value of a in \mathbf{A} and $a_{\mathbf{A} \dot{\cup} \neg U}$ of a in $\mathbf{A} \dot{\cup} \neg U$ (where U is the currently unfounded set), $a_{\mathbf{A} \wedge U}$ represents that a is true in \mathbf{A} and is contained in U , and $a_{\overline{\mathbf{A}} \vee U}$ represents that a is false in \mathbf{A} or it is contained in U .

To this end, a *set of assumptions* is a consistent set \mathcal{A} of signed literals. A solution \mathbf{A} to a nogood δ resp. a set Δ of nogoods *satisfies* \mathcal{A} , if $\mathcal{A} \subseteq \mathbf{A}$. That is, assumptions fix the truth value of some atoms. Modern ASP and SAT solvers support assumptions natively, and they can be easily undone without a complete reset of the reasoner and recreating the whole problem instance. This is an essential feature for efficiently implementing our improved encoding.

Our encoding Ω_{Π} is then

$$\Omega_{\Pi} = N_{\Omega, \Pi} \cup O_{\Omega, \Pi},$$

where $N_{\Omega, \Pi} = \{\{\mathbf{F}a \mid a \in A(\Pi)\}\} \cup \bigcup_{a \in A(\Pi)} D_a \cup \bigcup_{r \in \Pi} (H_r \cup C_r)$ is the necessary part and

- $\{\mathbf{F}a \mid a \in A(\Pi)\}$ encodes that we search for a nonempty unfounded set;

$$\bullet D_a = \begin{cases} \{\{\mathbf{F}a_{\mathbf{A} \wedge U}, \mathbf{T}a_{\mathbf{A}}, \mathbf{T}a\}, \{\mathbf{T}a_{\mathbf{A} \wedge U}, \mathbf{F}a_{\mathbf{A}}\}, \{\mathbf{T}a_{\mathbf{A} \wedge U}, \mathbf{F}a\}\} \cup \\ \{\{\mathbf{F}a_{\overline{\mathbf{A}} \vee U}, \mathbf{F}a_{\mathbf{A}}\}, \{\mathbf{F}a_{\overline{\mathbf{A}} \vee U}, \mathbf{T}a\}, \{\mathbf{T}a_{\overline{\mathbf{A}} \vee U}, \mathbf{T}a_{\mathbf{A}}, \mathbf{F}a\}\} \cup \\ \{\{\mathbf{T}a_{\mathbf{A} \dot{\cup} \neg U}, \mathbf{F}a_{\mathbf{A}}\}, \{\mathbf{T}a_{\mathbf{A} \dot{\cup} \neg U}, \mathbf{T}a\}, \{\mathbf{F}a_{\mathbf{A} \dot{\cup} \neg U}, \mathbf{T}a_{\mathbf{A}}, \mathbf{F}a\}\} \end{cases}$$

encodes that $a_{\mathbf{A} \wedge U}$ is true iff $a_{\mathbf{A}}$ and a are both true, $a_{\overline{\mathbf{A}} \vee U}$ is true iff $a_{\mathbf{A}}$ is false or a is true, and $a_{\mathbf{A} \dot{\cup} \neg U}$ is true iff $a_{\mathbf{A}}$ is true and a is false;

- $H_r = \{\{\mathbf{T}h_r\} \cup \{\mathbf{F}h \mid h \in H(r)\}\} \cup \{\{\mathbf{F}h_r, \mathbf{T}h\} \mid h \in H(r)\}$
encodes that h_r is true for a rule r iff some atom of $H(r)$ is in the unfounded set; and

$$\bullet C_r = \begin{cases} \{\{\mathbf{T}h_r\} \cup \\ \{\mathbf{T}a_{\mathbf{A}} \mid a \in B^+(\hat{r})\} \cup \{\mathbf{F}a_{\mathbf{A}} \mid a \in B^-(\hat{r})\} \cup & (i) \\ \{\mathbf{F}a_{\mathbf{A} \wedge U} \mid a \in B_o^+(r)\} \cup \{\mathbf{t}a \mid a \in B_e(\hat{r})\} \cup & (ii) \\ \{\mathbf{T}h_{\overline{\mathbf{A}} \vee U} \mid h \in H(r)\} & (iii) \end{cases}$$

encodes that if h_r is true, then one of (i), (ii) or (iii) in Definition 3.1 must hold.

More specifically, for an unfounded set U and a rule r with $H(r) \cap U \neq \emptyset$ (h_r is true) it must not happen that $\mathbf{A} \models B(r)$ (Condition (i) fails), no $a \in B_o^+(r)$ with $\mathbf{A} \models a$ is in the unfounded set and all $a \in B_e(\hat{r})$ are true under $\mathbf{A} \dot{\cup} \neg U$ (Condition (ii) fails), and all $h \in H(r)$ with $\mathbf{A} \models h$ are in the unfounded set (Condition (iii) fails).

Example 3.11. For $\Pi = \{r_1 : p \leftarrow \&id[p]()\}$ in Example 2.10, the constructed nogood set is

$$\begin{aligned} \Omega_{\Pi} = & \{\mathbf{F}p\}, \{\mathbf{F}p_{\mathbf{A} \wedge U}, \mathbf{T}p_{\mathbf{A}}, \mathbf{T}p\}, \{\mathbf{T}p_{\mathbf{A} \wedge U}, \mathbf{F}p_{\mathbf{A}}\}, \{\mathbf{T}p_{\mathbf{A} \wedge U}, \mathbf{F}p\}, \\ & \{\mathbf{F}p_{\overline{\mathbf{A}} \vee U}, \mathbf{F}p\}, \{\mathbf{F}p_{\overline{\mathbf{A}} \vee U}, \mathbf{T}p\}, \{\mathbf{T}p_{\overline{\mathbf{A}} \vee U}, \mathbf{T}p_{\mathbf{A}}, \mathbf{F}p\}, \{\mathbf{T}p_{\mathbf{A} \dot{\cup} \neg U}, \mathbf{F}p_{\mathbf{A}}\}, \\ & \{\mathbf{T}p_{\mathbf{A} \dot{\cup} \neg U}, \mathbf{T}p\}, \{\mathbf{F}p_{\mathbf{A} \dot{\cup} \neg U}, \mathbf{T}p_{\mathbf{A}}, \mathbf{F}p\}, \\ & \{\mathbf{T}h_{r_1}, \mathbf{F}p\}, \{\mathbf{F}h_{r_1}, \mathbf{T}p\}, \{\mathbf{T}h_{r_1}, \mathbf{T}e_{\&id[p]_{\mathbf{A}}}, \mathbf{T}e_{\&id[p]}, \mathbf{T}p_{\overline{\mathbf{A}} \vee U}\}. \end{aligned}$$

Towards computing unfounded sets, observe that every unfounded set can be extended to a solution to the set of nogoods Ω_{Π} over $A(\Omega_{\Pi})$. Conversely, the solutions to Ω_{Π} include specific extensions of all unfounded sets, which are again characterized by induced assignments; that is, by assigning true to all atoms in U , to all h_r such that $H(r)$ intersects with U , and to all replacement atoms $e_{\&g[p]}(\mathbf{c})$ such that $\&g[p](\mathbf{c})$ is true under $\mathbf{A} \dot{\cup} \neg U$, appropriate truth values to the auxiliary atoms according to their intuitive meaning, and assigning false to all other atoms in $A(\Omega_{\Pi})$. More formally, this leads us to the following assignment:

Definition 3.12 (Induced Assignment of an Unfounded Set wrt. Ω_Π). *Let U be an unfounded set of a program Π wrt. assignment \mathbf{A} . The assignment induced by U wrt. Ω_Π , denoted $I_\Omega(U, \Omega_\Pi, \Pi, \mathbf{A})$, is*

$$I_\Omega(U, \Omega_\Pi, \Pi, \mathbf{A}) = I_\Omega^0(U, \Pi, \mathbf{A}) \cup \{\mathbf{F}a \mid a \in A(\Omega_\Pi), \mathbf{T}a \notin I_\Omega^0(U, \Pi, \mathbf{A})\}$$

where

$$\begin{aligned} I_\Omega^0(U, \Pi, \mathbf{A}) = & \{\mathbf{T}a \mid a \in U\} \cup \{\mathbf{T}h_r \mid r \in \Pi, H(r) \cap U \neq \emptyset\} \cup \\ & \{\mathbf{T}e_{\&g[\mathbf{p}]}(\mathbf{c}) \mid e_{\&g[\mathbf{p}]}(\mathbf{c}) \in A(\hat{\Pi}), \mathbf{A} \dot{\cup} \neg.U \models \&g[\mathbf{p}](\mathbf{c})\} \cup \\ & \{\mathbf{T}a_{\mathbf{A}} \mid a \in A(\Pi), \mathbf{T}a \in \mathbf{A}\} \cup \{\mathbf{T}\hat{a}_{\mathbf{A}} \mid a \in EA(\Pi), \mathbf{A} \models a\} \cup \\ & \{\mathbf{T}a_{\mathbf{A} \wedge U} \mid a \in A(\Pi), \mathbf{T}a \in \mathbf{A}, a \in U\} \cup \\ & \{\mathbf{T}a_{\mathbf{A} \dot{\cup} \neg.U} \mid a \in A(\Pi), \mathbf{T}a \in \mathbf{A}, a \notin U\} \cup \\ & \{\mathbf{T}a_{\bar{\mathbf{A}} \vee U} \mid a \in A(\Pi), \mathbf{F}a \in \mathbf{A} \text{ or } a \in U\}. \end{aligned}$$

If we adopt for an assignment \mathbf{A} the assumption set

$$\begin{aligned} \mathcal{A}_{\mathbf{A}} = & \{\mathbf{T}a_{\mathbf{A}} \mid a \in A(\Pi), \mathbf{T}a \in \mathbf{A}\} \cup \{\mathbf{F}a_{\mathbf{A}} \mid a \in A(\Pi), \mathbf{F}a \in \mathbf{A}\} \cup \\ & \{\mathbf{T}\hat{a}_{\mathbf{A}} \mid a \in EA(\Pi), \mathbf{A} \models a\} \cup \{\mathbf{F}\hat{a}_{\mathbf{A}} \mid a \in EA(\Pi), \mathbf{A} \not\models a\}, \end{aligned}$$

then all assignments induced by unfounded sets of Π wrt. \mathbf{A} are solutions to Ω_Π wrt. $\mathcal{A}_{\mathbf{A}}$ (but not conversely, because intuitively the latter do not reflect the semantics of external sources).

As before, we call a set of nogoods N *conservative*, if $I_\Omega(U, \Omega_\Pi, \Pi, \mathbf{A})$ is a solution to N for every unfounded set U of Π wrt. \mathbf{A} . Under this property, these interpretations are solutions of the whole nogood set that comply with the assumptions from \mathbf{A} .

Proposition 3.13. *Let U be an unfounded set of a program Π wrt. assignment \mathbf{A} such that $\mathbf{A}^{\mathbf{T}} \cap U \neq \emptyset$. If $O_{\Omega, \Pi}$ is conservative, then $I_\Omega(U, \Omega_\Pi, \Pi, \mathbf{A})$ is a solution to Ω_Π that satisfies $\mathcal{A}_{\mathbf{A}}$.*

Corollary 3.14. *If Ω_Π has no solution which satisfies $\mathcal{A}_{\mathbf{A}}$, then $U \cap \mathbf{A}^{\mathbf{T}} = \emptyset$ for every unfounded set U of Π (assuming $O_{\Omega, \Pi}$ is conservative).*

The next property allows us to find the unfounded sets of Π wrt. \mathbf{A} among all solutions to $\Omega^{\mathbf{A}}$ that satisfy $\mathcal{A}_{\mathbf{A}}$ by using a postcheck on the external atoms.

Theorem 3.15. *Let S be a solution to Ω_Π (with conservative $O_{\Omega, \Pi}$) that satisfies $\mathcal{A}_{\mathbf{A}}$ such that*

(a) $\mathbf{T}e_{\&g[\mathbf{p}]}(\mathbf{c}) \in S$ and $\mathbf{A} \not\models \&g[\mathbf{p}](\mathbf{c})$ implies $\mathbf{A} \dot{\cup} \neg.U \models \&g[\mathbf{p}](\mathbf{c})$; and

(b) $\mathbf{F}e_{\&g[\mathbf{p}]}(\mathbf{c}) \in S$ and $\mathbf{A} \models \&g[\mathbf{p}](\mathbf{c})$ implies $\mathbf{A} \dot{\cup} \neg.U \not\models \&g[\mathbf{p}](\mathbf{c})$

where $U = \{ \mid a \in A(\Pi) \mid \mathbf{T}a \in S \}$. Then U is an unfounded set of Π wrt. \mathbf{A} .

As for $\Gamma_\Pi^{\mathbf{A}}$, the proposition states that the non-replacement atoms in S that are true and appear in Π form an unfounded set, provided that each replacement atom $e_{\&g[\mathbf{p}]}(\mathbf{c})$ in S has the same truth value as $\&g[\mathbf{p}](\mathbf{c})$ under $\mathbf{A} \dot{\cup} \neg.U$ (as in Definition 3.12). Again, this check is just required if the truth value of $e_{\&g[\mathbf{p}]}(\mathbf{c})$ in S is different from the one of $\&g[\mathbf{p}](\mathbf{c})$ under \mathbf{A} .

Example 3.16. *Reconsider program $\Pi = \{r_1 : p \leftarrow \&id[p]()\}$ from Example 2.10 and the compatible set $\mathbf{A}_2 = \{\mathbf{T}p, \mathbf{T}e_{\&id[p]}\}$. The nogood set*

$$\begin{aligned} \Omega_\Pi = & \{\mathbf{F}p\}, \{\mathbf{F}p_{\mathbf{A} \wedge U}, \mathbf{T}p_{\mathbf{A}}, \mathbf{T}p\}, \{\mathbf{T}p_{\mathbf{A} \wedge U}, \mathbf{F}p_{\mathbf{A}}\}, \{\mathbf{T}p_{\mathbf{A} \wedge U}, \mathbf{F}p\}, \\ & \{\mathbf{F}p_{\bar{\mathbf{A}} \vee U}, \mathbf{F}p\}, \{\mathbf{F}p_{\bar{\mathbf{A}} \vee U}, \mathbf{T}p\}, \{\mathbf{T}p_{\bar{\mathbf{A}} \vee U}, \mathbf{T}p_{\mathbf{A}}, \mathbf{F}p\}, \{\mathbf{T}p_{\mathbf{A} \dot{\cup} \neg.U}, \mathbf{F}p_{\mathbf{A}}\}, \\ & \{\mathbf{T}p_{\mathbf{A} \dot{\cup} \neg.U}, \mathbf{T}p\}, \{\mathbf{F}p_{\mathbf{A} \dot{\cup} \neg.U}, \mathbf{T}p_{\mathbf{A}}, \mathbf{F}p\}, \\ & \{\mathbf{T}h_{r_1}, \mathbf{F}p\}, \{\mathbf{F}h_{r_1}, \mathbf{T}p\}, \{\mathbf{T}h_{r_1}, \mathbf{T}e_{\&id[p]_{\mathbf{A}}}, \mathbf{T}e_{\&id[p]}, \mathbf{T}p_{\bar{\mathbf{A}} \vee U}\} \end{aligned}$$

with assumptions $\mathcal{A}_{\mathbf{A}_2} = \{\mathbf{T}p_{\mathbf{A}}\}$ has solutions $S \supseteq \{\mathbf{T}h_{r_1}, \mathbf{T}p, \mathbf{T}p_{\mathbf{A}}, \mathbf{F}e_{\&id[p]}, \mathbf{T}p_{\mathbf{A} \wedge U}, \mathbf{T}p_{\bar{\mathbf{A}} \vee U}, \mathbf{F}p_{\mathbf{A} \dot{\cup} \neg U}\}$, which correspond to the unfounded set $U = \{p\}$. Here, $\mathbf{F}e_{\&id[p]}()$ represents that $\mathbf{A}_2 \dot{\cup} \neg U \not\models \&id[p]()$.

We will see in Section 6 that the encoding Ω_{Π} is superior to $\Gamma_{\Pi}^{\mathbf{A}}$ for many practically relevant programs. The effect becomes especially visible if they need many unfounded set checks, which intuitively is the case when many answer sets exist; here reusability of the encoding is very beneficial, while for small programs with few answer sets, the incurred overhead does not lead to savings.

4 Optimization and Learning

In this section we first discuss some refinements and optimizations of our nogood encodings for UFS search. In particular, we present nogoods which prune irrelevant parts of the search space; they can be integrated into both encodings $\Gamma_{\Pi}^{\mathbf{A}}$ and Ω_{Π} under suitable adjustments. After that, we propose a strategy for learning nogoods from detected unfounded sets, avoiding that the same unfounded set is generated later again.

4.1 Optimization

We present now three optimizations which turned out to be effective in improving UFS search, where the second and the third exclude each other, i.e., they can not be used simultaneously.

O1: Restricting the UFS Search to atoms in the compatible set. First, not all atoms in a program are relevant for the unfounded set search: atoms that are false under \mathbf{A} can be ignored.

Proposition 4.1. *Suppose U is an unfounded set of Π wrt. a complete assignment \mathbf{A} such that $\mathbf{A} \not\models a$ for some $a \in U$. Then $U \setminus \{a\}$ is an unfounded set of Π wrt. \mathbf{A} .*

The nogoods for this optimization are simple. In the encoding $\Gamma_{\Pi}^{\mathbf{A}}$, we add the conservative nogood $\{\mathbf{T}a\}$ for each $a \in A(\Pi)$ with $\mathbf{A} \not\models a$ to the optimization part $O_{\Gamma, \Pi}^{\mathbf{A}}$ and in the encoding Ω_{Π} the conservative nogood $\{\mathbf{F}a_{\mathbf{A}}, \mathbf{T}a\}$ for each $a \in A(\Pi)$ to the optimization part $O_{\Omega, \Pi}$.

O2: Avoiding guesses of replacement atoms. In some situations, the truth value of a replacement atom b in a solution S to $\Gamma_{\Pi}^{\mathbf{A}}$ resp. Ω_{Π} with assumptions $\mathcal{A}_{\mathbf{A}}$ is irrelevant. That is, both $S_{\mathbf{T}b} = (S \setminus \{\mathbf{T}b, \mathbf{F}b\}) \cup \{\mathbf{T}b\}$ and $S_{\mathbf{F}b} = (S \setminus \{\mathbf{T}b, \mathbf{F}b\}) \cup \{\mathbf{F}b\}$ are solutions to $\Gamma_{\Pi}^{\mathbf{A}}$ resp. Ω_{Π} that satisfy $\mathcal{A}_{\mathbf{A}}$, and they represent the same unfounded set. We then can set the truth value of b to an (arbitrary) fixed value instead of inspecting both alternatives. The next proposition states a sufficient criterion for this irrelevance.

Proposition 4.2. *Let b be a replacement atom, and let S be a solution to $\Gamma_{\Pi}^{\mathbf{A}}$ resp. Ω_{Π} satisfying $\mathcal{A}_{\mathbf{A}}$. If for every rule $r \in \Pi$ such that $b \in B^+(\hat{r}) \cup B^-(\hat{r})$ and $\mathbf{A} \models B(r)$, either*

(a) *for some $a \in B_o^+(r)$ such that $\mathbf{A} \models a$, it holds that $\mathbf{T}a \in S$, or*

(b) *for some $a \in H(r)$ such that $\mathbf{A} \models a$, it holds that $\mathbf{F}a \in S$,*

then both $S_{\mathbf{T}b}$ and $S_{\mathbf{F}b}$ are solutions to $\Gamma_{\Pi}^{\mathbf{A}}$ resp. Ω_{Π} that satisfy $\mathcal{A}_{\mathbf{A}}$.

This property can be utilized by adding conservative nogoods. Recall that $A(\Gamma_{\Pi}^{\mathbf{A}})$ and $A(\Omega_{\Pi})$ contain atoms l_r for every $r \in \Pi$. They intuitively serve to encode for a solution S to $\Gamma_{\Pi}^{\mathbf{A}}$ resp. Ω_{Π} with assumptions $\mathcal{A}_{\mathbf{A}}$ whether the truth values of the replacement atoms in $B(r)$ are relevant or can be set arbitrarily. The

following nogoods label relevant rules r , forcing l_r to be false iff some of the conditions in Proposition 4.2 holds. For the encoding $\Gamma_{\Pi}^{\mathbf{A}}$, we add to $O_{\Gamma, \Pi}^{\mathbf{A}}$ for each rule r :

$$L_{\Gamma, r}^{\mathbf{A}} = \{\{\mathbf{T}l_r, \mathbf{T}a\} \mid a \in B_o^+(r), \mathbf{A} \models a\} \cup \{\{\mathbf{T}l_r, \mathbf{F}a\} \mid a \in H(r), \mathbf{A} \models a\} \cup \{\{\mathbf{F}l_r\} \cup \{\mathbf{F}a \mid a \in B_o^+(r), \mathbf{A} \models a\} \cup \{\mathbf{T}a \mid a \in H(r), \mathbf{A} \models a\}\}.$$

For the encoding Ω_{Π} , we add to $O_{\Omega, \Pi}$ for each rule r :

$$L_{\Omega, r} = \{\{\mathbf{T}l_r, \mathbf{T}a, \mathbf{T}a_{\mathbf{A}}\} \mid a \in B_o^+(r)\} \cup \{\{\mathbf{T}l_r, \mathbf{F}a, \mathbf{T}a_{\mathbf{A}}\} \mid a \in H(r)\} \cup \{\{\mathbf{F}l_r\} \cup \{\mathbf{F}a_{\mathbf{A} \wedge U} \mid a \in B_o^+(r)\} \cup \{\mathbf{T}a_{\bar{\mathbf{A}} \text{ or } U} \mid a \in H(r)\}\}.$$

These constraints exclusively enforce either $\mathbf{T}l_r$ or $\mathbf{F}l_r$. Hence, the truth value of l_r deterministically depends on the other atoms, i.e., the nogoods do not cause additional guessing.

By Proposition 4.2 we can set the truth value of a replacement atom b arbitrarily, if l_r is false for all r such that $b \in B^+(\hat{r})$ or $b \in B^-(\hat{r})$. As mentioned after Theorem 3.9, it is beneficial to set the truth value of $e_{\&g[p]}(\mathbf{c})$ to the one of $\&g[p](\mathbf{c})$ under \mathbf{A} , because this can reduce the number of external atoms that must be checked. The following nogoods enforce a coherent interpretation of the replacement atoms.

For the encoding $\Gamma_{\Pi}^{\mathbf{A}}$ we add to $O_{\Gamma, \Pi}^{\mathbf{A}}$ for each rule r :

$$F_{\Gamma, r}^{\mathbf{A}} = \{\{\mathbf{F}l_r \mid b \in B^+(\hat{r}) \cup B^-(\hat{r})\} \cup \{\mathbf{F}b\} \mid b \in B_e(\hat{r}), \mathbf{A} \models b\} \cup \{\{\mathbf{F}l_r \mid b \in B^+(\hat{r}) \cup B^-(\hat{r})\} \cup \{\mathbf{T}b\} \mid b \in B_e(\hat{r}), \mathbf{A} \not\models b\}$$

while for the encoding Ω_{Π} we add to $O_{\Omega, \Pi}$ for each rule r :

$$F_{\Omega, r} = \{\{\mathbf{F}l_r \mid b \in B^+(\hat{r}) \cup B^-(\hat{r})\} \cup \{\mathbf{T}b_{\mathbf{A}}, \mathbf{F}b\} \mid b \in B_e(\hat{r})\} \cup \{\{\mathbf{F}l_r \mid b \in B^+(\hat{r}) \cup B^-(\hat{r})\} \cup \{\mathbf{F}b_{\mathbf{A}}, \mathbf{T}b\} \mid b \in B_e(\hat{r})\}$$

In summary, the encoding $\Gamma_{\Pi}^{\mathbf{A}}$ has the optimization part $O_{\Gamma, \Pi}^{\mathbf{A}} = \bigcup_{r \in \Pi} L_{\Gamma, r}^{\mathbf{A}} \cup F_{\Gamma, r}^{\mathbf{A}}$ and the encoding Ω_{Π} the optimization part $O_{\Omega, \Pi} = \bigcup_{r \in \Pi} L_{\Omega, r} \cup F_{\Omega, r}$.

We give now an example for this optimization using our encoding Γ .

Example 4.3. Consider the program $\Pi = \{r_1: p \leftarrow \&id[p](); r_2: q \leftarrow \&id[q]()\}$, and the compatible set $\hat{\mathbf{A}} = \{\mathbf{T}p, \mathbf{F}q, \mathbf{T}e_{\&id[p]}(), \mathbf{F}e_{\&id[q]}()\}$. Then the necessary part of encoding $\Gamma_{\Pi}^{\mathbf{A}}$ has solutions $S_1 \supseteq \{\mathbf{T}h_{r_1}, \mathbf{T}p, \mathbf{F}e_{\&id[p]}(), \mathbf{F}h_{r_2}, \mathbf{F}q, \mathbf{F}e_{\&id[q]}()\}$ and $S_2 \supseteq \{\mathbf{T}h_{r_1}, \mathbf{T}p, \mathbf{F}e_{\&id[p]}(), \mathbf{F}h_{r_2}, \mathbf{F}q, \mathbf{T}e_{\&id[q]}()\}$ (which represent the same unfounded set $U = \{p\}$). Here, the optimization part for r_2 , $L_{r_2}^{\mathbf{A}} \cup F_{r_2}^{\mathbf{A}} = \{\{\mathbf{T}l_{r_2}, \mathbf{F}q\}, \{\mathbf{F}l_{r_2}, \mathbf{T}q\}, \{\mathbf{F}l_{r_2}, \mathbf{T}e_{\&id[q]}()\}\}$, eliminates solutions S_2 for $\Gamma_{\Pi}^{\mathbf{A}}$. This is beneficial as for solutions S_1 the postcheck is easier ($e_{\&id[q]}()$ in S_1 and $\&id[q]()$ have the same truth value under \mathbf{A}).

O3: Exchanging nogoods between UFS and main search. The third optimization allows for the exchange of learned knowledge about external atoms between the UFS check and the main search for compatible sets. For this purpose, we first define nogoods which correctly describe the input-output relationship of external atoms.

Definition 4.4. A nogood of the form $N = \{\mathbf{T}t_1, \dots, \mathbf{T}t_n, \mathbf{F}f_1, \dots, \mathbf{F}f_m, \sigma e_{\&g[p]}(\mathbf{c})\}$, where σ is \mathbf{T} or \mathbf{F} , is a valid input-output-relationship, if for every assignment \mathbf{A} such that $N \setminus \{\sigma e_{\&g[p]}(\mathbf{c})\} \subseteq \mathbf{A}$ it holds that $\mathbf{A} \models \&g[p](\mathbf{c})$ if $\sigma = \mathbf{F}$, and $\mathbf{A} \not\models \&g[p](\mathbf{c})$ if $\sigma = \mathbf{T}$.

Let N be a nogood which is a valid input-output-relationship learned during the *main search*, i.e., for compatible sets of $\hat{\Pi}$, and let $\bar{\mathbf{F}} = \mathbf{T}$ and $\bar{\mathbf{T}} = \mathbf{F}$.

Definition 4.5 (Nogood Transformation \mathcal{T}_Γ). *For a valid input-output relationship $N = \{\mathbf{T}t_1, \dots, \mathbf{T}t_n, \mathbf{F}f_1, \dots, \mathbf{F}f_m, \sigma e_{\&g[p]}(\mathbf{c})\}$ and an assignment \mathbf{A} , the nogood transformation \mathcal{T}_Γ is defined as*

$$\mathcal{T}_\Gamma(N, \mathbf{A}) = \begin{cases} \emptyset & \text{if } \mathbf{F}t_i \in \mathbf{A} \text{ for some } 1 \leq i \leq n, \\ \{ \{ \mathbf{F}t_1, \dots, \mathbf{F}t_n \} \cup \{ \sigma e_{\&g[p]}(\mathbf{c}) \} \} \cup \\ \{ \{ \mathbf{T}f_i \mid 1 \leq i \leq m, \mathbf{A} \models f_i \} \} & \text{otherwise.} \end{cases}$$

The next result states that $\mathcal{T}_\Gamma(N, \mathbf{A})$ can be considered, for all valid input-output relationships N under all assignments \mathbf{A} , without losing unfounded sets.

Proposition 4.6. *Let N be a valid input-output relationship, and let U be an unfounded set wrt. Π and \mathbf{A} . If $O_{\Gamma, \Pi}^{\mathbf{A}}$ contains only conservative nogoods, then $I_\Gamma(U, \Gamma_{\Pi}^{\mathbf{A}}, \Pi, \mathbf{A})$ is a solution to $\mathcal{T}_\Gamma(N, \mathbf{A})$ (i.e., also nogoods $\mathcal{T}_\Gamma(N, \mathbf{A})$ are conservative).*

Hence, all valid input-output relationships for external atoms that are learned during the search for compatible sets can be reused (applying the above transformation) for the UFS check. Moreover, during the evaluation of external atoms in the postcheck for candidate unfounded sets (i.e., solutions to $\Gamma_{\Pi}^{\mathbf{A}}$), further valid input-output relationships might be learned. They can in turn be used by (further) unfounded set checks (in transformed form) but also in the search for compatible sets.

Example 4.7 (Set Partitioning). *For the program Π from Example 2.5, consider the compatible set $\hat{\mathbf{A}}^{\mathbf{T}} = \{ \text{domain}(a), \text{sel}(a), e_{\&diff[n\text{sel}]}(a) \}$. Suppose the main search has learned the input-output relationship $N = \{ \mathbf{T} \text{domain}(a), \mathbf{F} \text{n\text{sel}}(a), \mathbf{F} e_{\&diff[n\text{sel}]}(a) \}$. Then the transformed nogood is $a\mathcal{T}_\Gamma(N, \mathbf{A}) = \{ \{ \mathbf{F} \text{domain}(a), \mathbf{F} e_{\&diff[n\text{sel}]}(a) \} \}$; it intuitively encodes that if $\text{domain}(a)$ is not in the unfounded set U , then $e_{\&diff[n\text{sel}]}(a)$ is true under $\mathbf{A} \dot{\cup} \neg.U$. This holds because $e_{\&diff[n\text{sel}]}(a)$ is true under \mathbf{A} and can only change its truth value if $\text{domain}(a)$ becomes false.*

This learning technique can be adopted for our encoding Ω_Π as follows.

Definition 4.8 (Nogood Transformation \mathcal{T}_Ω). *For a valid input-output relationship N , the nogood transformation \mathcal{T}_Ω is defined as*

$$\mathcal{T}_\Omega(N) = \{ \{ \sigma e_{\&g[p]}(\mathbf{c}) \} \} \cup \{ \{ \mathbf{T}t_{1\mathbf{A}}, \mathbf{F}t_1, \dots, \mathbf{T}t_{n\mathbf{A}}, \mathbf{F}t_n, \mathbf{F}f_{1_{\mathbf{A} \dot{\cup} \neg.U}}, \dots, \mathbf{F}f_{m_{\mathbf{A} \dot{\cup} \neg.U}} \} \}.$$

Compared to $\mathcal{T}_\Gamma(N, \mathbf{A})$, the main difference is that $\mathcal{T}_\Omega(N)$ is independent of the current assignment, similar to the definition of our unfounded set detection problem Ω_Π .

The next result states that $\mathcal{T}_\Omega(N)$ can be considered, for all valid input-output relationships N under all assignments \mathbf{A} , without losing unfounded sets.

Proposition 4.9. *Let N be a valid input-output relationship, and let U be an unfounded set wrt. Π and \mathbf{A} . If $O_{\Omega, \Pi}$ contains only conservative nogoods, then $I_\Omega(U, \Omega_\Pi, \Pi, \mathbf{A})$ is a solution to $\mathcal{T}_\Omega(N)$ (i.e., also nogoods $\mathcal{T}_\Omega(N, \mathbf{A})$ are conservative).*

Hence, all valid input-output relationships for external atoms that are learned during the search for compatible sets can be reused (applying the above transformation) for the unfounded set check. Moreover, during the evaluation of external atoms in the postcheck for candidate unfounded sets (i.e., solutions to Ω_Π), further valid input-output relationships might be learned. The latter can be used in turn in the main search for compatible sets and (in transformed form) in further unfounded set checks.

Example 4.10 (cont'd). *Reconsider the program Π from Example 2.5 and the compatible set $\mathbf{A}^T = \{domain(a), sel(a), e_{\&diff}[domain, nsel](a)\}$. Suppose the main search has learned the input-output relationship $N = \{\mathbf{T}domain(a), \mathbf{F}nsel(a), \mathbf{F}e_{\&diff}[domain, nsel](a)\}$. The transformed nogood is*

$$\mathcal{T}_\Omega(N) = \{\{\mathbf{T}domain(a)_{\mathbf{A}}, \mathbf{F}domain(a), \mathbf{F}nsel(a)_{\mathbf{A} \dot{\cup} \neg.U}, \mathbf{F}e_{\&diff}[domain, nsel](a)_{\mathbf{A} \dot{\cup} \neg.U}\}\};$$

it intuitively encodes that if $domain(a)$ is true in the assignment \mathbf{A} but not in the unfounded set U , and if $nsel(a)$ is false in $\mathbf{A} \dot{\cup} \neg.U$, then $e_{\&diff}[domain, nsel](a)$ is true under $\mathbf{A} \dot{\cup} \neg.U$. This holds as $e_{\&diff}[domain, nsel](a)$ is true under \mathbf{A} and can only change its truth value if $domain(a)$ gets false.

The nogood exchange also benefits from our assignment-independent encoding. With our previous encoding $\Gamma_{\Pi}^{\mathbf{A}}$, we needed to build the SAT instance from scratch for every unfounded set check. Thus, nogoods learned in the main search for compatible sets need to be transformed and added to the UFS detection problem for every check (otherwise they are lost). With our new encoding Ω_{Π} , this is done only once because learned nogoods are kept for multiple unfounded set checks. This also allows us to make use of advanced forgetting heuristics in SAT solvers more effective.

Finally, an important note is that the optimizations O2 and O3 *can not* be used simultaneously (differently from O1 and O2 resp. O1 and O3), as this can result in contradictions due to (transformed) learned nogoods. That is, O2 and O3 together are not conservative (we thus disabled O2 in our experiments).

4.2 Learning Nogoods from Unfounded Sets

Until now we have considered merely detecting unfounded sets. A strategy to *learn* from detected unfounded sets for the main search for compatible sets is missing. We next develop such a strategy which we call *unfounded set learning* (UFL).

Example 4.11. *Consider the program $\Pi = \{p \leftarrow \&id[p](); x_1 \vee x_2 \vee \dots \vee x_k \leftarrow\}$. As we know from Example 3.4, $U = \{p\}$ is a UFS of the subprogram $\Pi' = \{p \leftarrow \&id[p]()\}$ wrt. $\mathbf{A} = \{\mathbf{T}p, \mathbf{T}e_{\&id}()\}$. The same is true for Π and moreover for every $\mathbf{A}' \supset \mathbf{A}$; i.e., p must never be true.*

The program in Example 4.11 has many compatible sets, and half of them (all where p is true) will fail the UFS check for the same reason. We thus develop a strategy for generating additional nogoods to guide the search for compatible sets in a way such that the same unfounded sets are not reconsidered. We present two such strategies, but will focus on the first one because our experiments have shown that the first one is superior for all our instances.

4.2.1 UFS-based Learning

For an unfounded set U of Π wrt. \mathbf{A} we define a set $L_1(U, \Pi, \mathbf{A})$ of learned nogoods as follows. Suppose that r_1, \dots, r_j are all rules r in Π such that $H(r) \cap U \neq \emptyset$ and $U \cap B_o^+(r) = \emptyset$, i.e., the set of all “external” rules of Π wrt. U (rules which do not directly depend positively on U). Then

$$L_1(U, \Pi, \mathbf{A}) = \{\{\sigma_0, \sigma_1, \dots, \sigma_j\} \mid \sigma_0 \in \{\mathbf{T}a \mid a \in U\}, \sigma_i \in H_i \text{ for all } 1 \leq i \leq j\},$$

where $H_i = \{\mathbf{T}h \mid h \in H(r_i) \setminus U, \mathbf{A} \models h\} \cup \{\mathbf{F}b \mid b \in B_o^+(r_i), \mathbf{A} \not\models b\}$. Formally we can show that adding this set of nogoods is correct, i.e., does not prune answer sets:

Proposition 4.12. *If U is an unfounded set of Π wrt. \mathbf{A} , then every answer set of Π is a solution to the nogoods in $L_1(U, \Pi, \mathbf{A})$.*

Example 4.13. Consider the program Π from Example 4.11 and suppose we have found the unfounded set $U = \{p\}$ wrt. the interpretation $\mathbf{A} = \{\mathbf{T}p, \mathbf{T}x_1\} \cup \{\mathbf{F}a_i \mid 1 < i \leq k\}$. Then the learned nogood $L_1(U, \mathbf{A}, \Pi) = \{\mathbf{T}p\}$ immediately guides the search to the part of the search tree where p is false, i.e., roughly half of the guesses are avoided.

4.2.2 Reduct-based Learning

A different learning strategy is based on the models of $f\Pi^{\mathbf{A}}$ rather than the unfounded set U itself, hinging on the observation that for every unfounded set U , $\mathbf{A} \dot{\cup} \neg.U$ is a model of $f\Pi^{\mathbf{A}}$; hence $U \neq \emptyset$ refutes \mathbf{A} as a minimal model of $f\Pi^{\mathbf{A}}$. This was noted in [Faber et al., 2011] for aggregates.

We exploit this to construct nogoods from a nonempty U wrt. a model \mathbf{A} as follows. The interpretation $\mathbf{A} \dot{\cup} \neg.U$ is not only a model of $f\Pi^{\mathbf{A}}$, but of *all* programs $\Pi' \subseteq f\Pi^{\mathbf{A}}$. Hence, if an assignment \mathbf{A}' falsifies *at least* the rules of Π which \mathbf{A} falsifies, and $\mathbf{A}'^{\mathbf{T}} \supset (\mathbf{A} \dot{\cup} \neg.U)^{\mathbf{T}}$, then \mathbf{A}' is not an answer set of Π . This yields the following nogood set $L_2(U, \Pi, \mathbf{A})$. Suppose r_1, \dots, r_n are all rules r of Π which are *not* in its FLP-reduct wrt. \mathbf{A} (i.e., $\mathbf{A} \not\models B(r_i)$). Then

$$L_2(U, \Pi, \mathbf{A}) = \{\{\mathbf{T}a \mid a \in (\mathbf{A} \dot{\cup} \neg.U)^{\mathbf{T}}\} \cup \{\sigma_0, \sigma_1, \dots, \sigma_n\} \\ \mid \sigma_0 \in \{\mathbf{T}a \mid a \in U\}, \sigma_i \in H_i \text{ for all } 1 \leq i \leq n\}$$

where $H_i = \{\mathbf{t}a \mid a \in B(\hat{r}), \hat{\mathbf{A}} \not\models a\}$, $1 \leq i \leq n$. That is, each nogood consists of the true-part of the smaller model $\mathbf{A} \dot{\cup} \neg.U$ of the reduct $f\Pi^{\mathbf{A}}$, an unfounded atom σ_0 (i.e. true in \mathbf{A} but not in $\mathbf{A} \dot{\cup} \neg.U$), and a false body literal σ_i ($1 \leq i \leq n$) for each rule of Π with unsatisfied body wrt. \mathbf{A} .

Example 4.14. Let $\Pi = \{p \leftarrow \&id[p](); q \leftarrow \&id[q]()\}$, where $\&id[a]()$ evaluates to true iff a is true. Suppose $\mathbf{A} = \{p, q\}$. Then $U = \{p, q\}$ is an unfounded set wrt. \mathbf{A} . In the above construction rule we have $\mathbf{A} \dot{\cup} \neg.U = \{\}$, $\sigma_0 \in \{\mathbf{T}p, \mathbf{T}q\}$ and $n = 0$ (because both rule bodies are satisfied wrt. \mathbf{A}). The learned nogoods are $L_2(U, \Pi, \mathbf{A}) = \{\{\mathbf{T}p\}, \{\mathbf{T}q\}\}$.

In Example 4.14, the learned nogoods will immediately guide the search to the interpretation $\{\mathbf{F}p, \mathbf{F}q\}$, which is the only one which becomes an answer set. Formally, we can show:

Proposition 4.15. If U is an unfounded set of Π wrt. \mathbf{A} and $\mathbf{A} \models \Pi$, then each answer set of Π is a solution to all nogoods in $L_2(U, \Pi, \mathbf{A})$.

However, $L_2(U, \Pi, \mathbf{A})$ appeared to be clearly inferior to $L_1(U, \Pi, \mathbf{A})$ from Section 4.2.1. Informally, its nogoods overfit the detected unfounded set and do not generalize well to other ones.

5 Deciding the Necessity of the Minimality Check

Although the minimality check based on unfounded sets is more efficient than the explicit minimality check, the computational costs are still high. Moreover, during the evaluation of $\hat{\Pi}$ for computing the compatible set $\hat{\mathbf{A}}$, the ordinary ASP solver has already made an unfounded set check, and we can safely assume that it is founded from the perspective of the ordinary ASP solver. Hence, all remaining unfounded sets which were not discovered by the ordinary ASP solver must involve external sources, as their behavior is not fully captured by the ASP solver.

In this section we pursue these ideas and give a decision criterion for deciding whether a further UFS check is necessary. We eventually define a class of programs which needs no additional UFS check. Intuitively,

we show that every unfounded set that is not already detected during the construction of $\hat{\mathbf{A}}$ contains input atoms of external atoms involved in cycles. If the program has no such input atom, then the UFS check is superfluous. Afterwards, we show how to apply this criterion, which holds in practically relevant cases, to program components; this often yields additional speedup. However, there are also cases where the UFS check can not be skipped; e.g., recursive URL retrieval from a web resource (which requires cyclic use of an external atom).

5.1 Basic Decision Criterion

We start with a definition of atom dependency.

Definition 5.1 (Atom Dependency). *For a ground program Π , and ground atoms $p(\mathbf{c})$ and $q(\mathbf{d})$, we say that*

- (i) $p(\mathbf{c})$ depends on $q(\mathbf{d})$, denoted $p(\mathbf{c}) \rightarrow q(\mathbf{d})$, if for some rule $r \in \Pi$ we have $p(\mathbf{c}) \in H(r)$ and $q(\mathbf{d}) \in B(r)$;
- (ii) $p(\mathbf{c})$ depends externally on $q(\mathbf{d})$, denoted $p(\mathbf{c}) \rightarrow_e q(\mathbf{d})$, if some rule $r \in \Pi$ and external atom $\&g[q_1, \dots, q_n](\mathbf{e}) \in B^+(r) \cup B^-(r)$ exist such that $p(\mathbf{c}) \in H(r)$ and $q \in \{q_1, \dots, q_n\}$.

In the following, we consider *dependency graphs* $G_{\Pi}^R = (V, E)$ for a ground program Π , whose vertices V are the ground atoms and whose edges E are given by a binary relation R over ground atoms ($E = R$). We omit R if $R = \rightarrow \cup \rightarrow_e$; we call $p(\mathbf{c}) \rightarrow q(\mathbf{d})$ also an *ordinary edge* and $p(\mathbf{c}) \rightarrow_e q(\mathbf{d})$ an *e-edge*.

We establish a lemma that allows us to restrict our attention to the “core” of an unfounded set, i.e., its most essential part; we can disregard atoms in a cut, which is defined as follows.

Definition 5.2 (Cut). *Let U be an unfounded set of Π wrt. \mathbf{A} . A set of atoms $C \subseteq U$ is a cut, if*

- (i) $b \not\rightarrow_e a$, for all $a \in C$ and $b \in U$ (C has no incoming e-edge from U), and
- (ii) $b \not\rightarrow a$ and $a \not\rightarrow b$, for all $a \in C$ and $b \in U \setminus C$ (there are no ordinary edges between C and $U \setminus C$).

We first prove that cuts can be removed from unfounded sets and the resulting set is still an unfounded set.

Lemma 5.3 (Unfounded Set Reduction Lemma). *Let U be an unfounded set of Π wrt. a complete assignment \mathbf{A} , and let C be a cut. Then, $Y = U \setminus C$ is an unfounded set of Π wrt. \mathbf{A} .*

Example 5.4. *Consider the following program:*

$$\Pi = \{r \leftarrow \&id[r](); \quad p \leftarrow \&id[r](); \quad p \leftarrow q; \quad q \leftarrow p\}.$$

Then we have $p \rightarrow q$, $q \rightarrow p$, $r \rightarrow_e r$ and $p \rightarrow_e r$. Π has the unfounded set $U = \{p, q, r\}$ wrt. $\mathbf{A} = \{\mathbf{T}p, \mathbf{T}q, \mathbf{T}r\}$. Observe that $C = \{p, q\}$ is a cut, and therefore $U \setminus C = \{r\}$ is an unfounded set of Π wrt. \mathbf{A} .

Next we prove that intuitively, for each unfounded set U of Π , either the input to some external atom is unfounded itself, or U is already detected when $\hat{\Pi}$ is evaluated.

Lemma 5.5 (EA-Input Unfoundedness). *Let U be an unfounded set of Π wrt. an assignment \mathbf{A} . If no edge $x \rightarrow_e y$ with $x, y \in U$ exists, then U is an unfounded set of $\hat{\Pi}$ wrt. $\hat{\mathbf{A}}$.*

Example 5.6. Reconsider the program Π from Example 5.4. Then the unfounded set $U' = \{p, q\}$ wrt. $\mathbf{A}' = \{\mathbf{T}p, \mathbf{T}q, \mathbf{F}r\}$ is already detected when

$$\hat{\Pi} = \{e_{\&id[r]}() \vee ne_{\&id[r]}() \leftarrow ; \quad r \leftarrow e_{\&id[r]}(); \quad p \leftarrow e_{\&id[r]}(); \quad p \leftarrow q; \quad q \leftarrow p\}$$

is evaluated by the ordinary ASP solver because $p \not\rightarrow_e q$ and $q \not\rightarrow_e p$. In contrast, the unfounded set $U'' = \{p, q, r\}$ wrt. $\mathbf{A}'' = \{\mathbf{T}p, \mathbf{T}q, \mathbf{T}r\}$ is not detected by the ordinary ASP solver because $p, r \in U''$ and $p \rightarrow_e r$.

Thus, the unfounded sets of Π wrt. \mathbf{A} that are not recognized during the evaluation of $\hat{\Pi}$ have cyclic dependencies over input atoms of some external atom. Programs with acyclic dependencies do not need additional UFS checks.

Recall that a *cycle* wrt. a binary relation R is a sequence $C = c_0, c_1, \dots, c_n, c_{n+1}$ of elements, $n \geq 0$, such that $(c_i, c_{i+1}) \in R$ for all $0 \leq i \leq n$ and $c_0 = c_{n+1}$. A set S contains a cycle under R , if there is a cycle $C = c_0, c_1, \dots, c_n, c_{n+1}$ under R such that $c_i \in S$ for all $0 \leq i \leq n + 1$.

Informally, the next proposition states that each unfounded set of Π wrt. \mathbf{A} which contains no cycle through the input atoms to some external atom corresponds to some unfounded set of $\hat{\Pi}$ wrt. $\hat{\mathbf{A}}$, i.e., the unfoundedness is already detected when $\hat{\Pi}$ is evaluated.

Let $\rightarrow^d = \rightarrow \cup \leftarrow \cup \rightarrow_e$, where \leftarrow is the inverse of \rightarrow (i.e. $\leftarrow = \{(x, y) \mid (y, x) \in \rightarrow\}$). A cycle $c_0, c_1, \dots, c_n, c_{n+1}$ under \rightarrow^d is called an *e-cycle*, if it contains an e-edge, i.e., $c_i \rightarrow_e c_{i+1}$ for some $0 \leq i \leq n$.

Theorem 5.7 (Relevance of e-cycles). *Let $U \neq \emptyset$ be an unfounded set of Π wrt. an interpretation \mathbf{A} which does not contain any e-cycle under \rightarrow^d . Then $\hat{\Pi}$ has a nonempty unfounded set wrt. $\hat{\mathbf{A}}$.*

Corollary 5.8. *If a program Π has no e-cycle under \rightarrow^d and $\hat{\Pi}$ has no unfounded set wrt. an interpretation $\hat{\mathbf{A}}$, then \mathbf{A} is unfounded-free for Π .*

This corollary can be used to increase the performance of an evaluation algorithm as follows: if there is no cycle under \rightarrow^d containing e-edges, then an explicit unfounded set check is not necessary because the unfounded set check during the evaluation of $\hat{\Pi}$ is sufficient. Note that this test can be done efficiently (in fact in linear time, similar to deciding stratifiability of an ordinary logic program). Moreover, in practice one can abstract from \rightarrow^d by using analogous relations on the level of predicate symbols instead of atoms. Clearly, if there is no e-cycle in the predicate dependency graph, then there can also be no e-cycle in the atom dependency graph. Hence, the predicate dependency graph can be safely used to decide whether the unfounded set check can be skipped.

Example 5.9. *All example programs so far need an UFS check, but the program $\Pi = \{out(X) \leftarrow \&diff[set_1, set_2](X)\} \cup F$, where $diff$ computes the set difference of unary predicates set_1 and set_2 and F is any set of facts, needs no UFS check as there is no e-cycle under \rightarrow^d . Also the program $\Pi = \{str(Z) \leftarrow dom(Z), str(X), str(Y), not \&concat[X, Y](Z)\}$ (where $\&concat$ takes two constants and computes their string concatenation) needs no UFS check; there is a cycle over an external atom, but no e-cycle under \rightarrow^d .*

Unfortunately, the converse of Theorem 5.7 does not hold, that is, $\hat{\Pi}$ may fail to be unfounded-free wrt. $\hat{\mathbf{A}}$ but no unfounded set of Π wrt. \mathbf{A} contains an e-cycle; thus, the condition in Corollary 5.8 is not necessary for unfounded-freeness of Π wrt. \mathbf{A} . However, the following generalization of Theorem 5.7 allows us to conclude that if $\hat{\Pi}$ is unfounded-free wrt. $\hat{\mathbf{A}}$, then every unfounded set U of Π wrt. \mathbf{A} must contain an atom that provides input to an external atom on a cycle under \rightarrow^d .

Definition 5.10 (Cyclic Input Atoms). *For a program Π , an atom a is a cyclic input atom, if some edge $b \rightarrow_e a$ with a path from a to b under \rightarrow^d exists.*

Let $CA(\Pi)$ denote the set of all cyclic input atoms of program Π .

Theorem 5.11 (Unfoundedness of Cyclic Input Atom). *Let $U \neq \emptyset$ be an unfounded set of Π wrt. \mathbf{A} such that $U \cap CA(\Pi) = \emptyset$. Then, $\hat{\Pi}$ has a nonempty unfounded set wrt. $\hat{\mathbf{A}}$.*

As a consequence of Theorem 5.11, we can add the nogood $\{\mathbf{F}a \mid a \in CA(\Pi)\}$ to $\Gamma_{\Pi}^{\mathbf{A}}$. Again using predicate symbols instead of atoms reduces the overhead of the dependency graph.

Example 5.12. *Reconsider Π in Example 5.4. Then $U = \{p, q\}$ is an unfounded set wrt. $\mathbf{A} = \{\mathbf{T}p, \mathbf{T}q, \mathbf{F}r\}$; as U is disjoint from $CA(\Pi) = \{r\}$, it is detected during the evaluation of $\hat{\Pi}$.*

5.2 Program Decomposition

The usefulness of the decision criterion can be increased by decomposing the program into components, such that the criterion can be applied componentwise. This allows us to restrict the UFS check to components with e-cycles, while e-cycle-free components can be ignored.

Let \mathcal{C} be a partitioning of the ordinary atoms $A(\Pi)$ of Π into subset-maximal strongly connected components under $\rightarrow \cup \rightarrow_e$. We define for each partition $C \in \mathcal{C}$ the subprogram Π_C associated with C as $\Pi_C = \{r \in \Pi \mid H(r) \cap C \neq \emptyset\}$.

We next show that if a program has an unfounded set U wrt. \mathbf{A} , then $U \cap C$ is an unfounded set wrt. \mathbf{A} for the subprogram of some strongly connected component C .

Theorem 5.13. *Let $U \neq \emptyset$ be an unfounded set of Π wrt. \mathbf{A} . Then, for some $C \in \mathcal{C}$ it holds that $U \cap C$ is a nonempty unfounded set of Π_C wrt. \mathbf{A} .*

Note that constraints (i.e., rules with empty head) do not harm this proposition. Each constraint r of kind $\leftarrow B(r)$ can be rewritten to $p \leftarrow B(r), \text{not } p$ for a new atom p , and $C = \{p\}$ is a strongly connected component with $\Pi_C = \{r\}$, which does not contain an e-cycle. Thus, for the rewritten constraints the according subprograms Π_C can be ignored anyways.

This proposition states that a search for unfounded sets can be done independently for the subprograms Π_C for all $C \in \mathcal{C}$. If a global unfounded set exists, then also for some program component exists. We know by Corollary 5.8 that programs Π without e-cycles can only contain unfounded sets that are already detected when $\hat{\Pi}$ is solved. If we apply Theorem 5.13 to the subprograms Π_C , we can safely ignore e-cycle-free program components.

Example 5.14. *Reconsider the program Π from Example 5.4. Then \mathcal{C} contains the components $C_1 = \{p, q\}$ and $C_2 = \{r\}$ and we have $\Pi_{C_1} = \{p \leftarrow \&id[r](); p \leftarrow q; q \leftarrow p\}$ and $\Pi_{C_2} = \{r \leftarrow \&id[r]()\}$. By Theorem 5.13, each unfounded set of Π wrt. some assignment \mathbf{A} gives rise to an unfounded set of either Π_{C_1} or Π_{C_2} . E.g., consider $U = \{p, q, r\}$ and $\mathbf{A} = \{\mathbf{T}p, \mathbf{T}q, \mathbf{T}r\}$; then $U \cap \{r\} = \{r\}$ is an unfounded set of Π_{C_2} wrt. \mathbf{A} . As Π_{C_1} has no e-cycles, we conclude from Corollary 5.8 that all unfounded sets of Π_{C_1} are already detected when $\hat{\Pi}$ (resp., $\hat{\Pi}_{C_1}$) is evaluated. Hence, only Π_{C_2} needs an additional UFS. Indeed, the only unfounded set of $\hat{\Pi}$ that is not detected when $\hat{\Pi}$ is evaluated is $\{r\}$, which is unfounded wrt. each interpretation $\mathbf{A} \supseteq \{\mathbf{T}r\}$ for Π_{C_2} and Π .*

Finally, we show that splitting, i.e., the component-wise check for foundedness, does not lead to spurious unfounded sets.

Proposition 5.15. *If U is an unfounded set of Π_C wrt. \mathbf{A} such that $U \subseteq C$, then U is an unfounded set of Π wrt. \mathbf{A} .*

The results can be generalized to subprograms that are larger than strongly connected components; however, we leave a detailed study of this for future work.

6 Implementation and Evaluation

For implementing our technique, we integrated CLASP into our prototype system dlvhex; we use CLASP as an ASP solver for computing compatible sets and as a SAT solver for solving the nogood set of the UFS check.

Regarding a concrete setting, there is a large number of combinations of the techniques in [Eiter et al., 2012a] and from above. Indeed, we may either activate or deactivate external behavior learning (EBL) and use either the explicit or the UFS-based minimality check. In the latter case, we can further use unfounded set learning (UFL), the decision criterion for skipping the unfounded set check can be exploited or ignored, and program decomposition might be used or not. Moreover, we can choose between the encodings Γ and Ω . In total, these are 34 different settings.

However, we will restrict our discussion to some interesting configurations. In general, we will activate the developed features stepwise such that in our tables the efficiency increases from left to right. We will start with the traditional algorithm based on an explicit minimality check without any kind of learning. In the next step we will add external behavior learning (EBL) [Eiter et al., 2012a], while UFL is not possible with the explicit check. Then we switch from the explicit minimality check to the UFS-based check without learning and without exploiting the decision criterion and program decomposition. Nevertheless, this naive kind of UFS-based minimality checking is usually already more efficient than the explicit minimality check with EBL. In the next step, we add the decision criterion and program decomposition. In the following, *monolithic (mol)* means that both the decision criterion and the program decomposition are off, and *modular (mod)* that they are on. Next we add EBL and UFL to the UFS-based minimality check, and finally we switch the encoding from Γ to Ω . However, we might skip some of the steps for specific benchmarks and argue why they are uninteresting in the respective cases. Detailed instance information and results with all combinations of parameters are available at <http://www.kr.tuwien.ac.at/research/projects/hexhex/ufs>.

Briefly, our results show a clear improvement, for both synthetic and application instances, by the UFS check and EBL. Moreover, a closer analysis shows that the UFS check decreases in some cases not only the runtime but also the number of enumerated candidates (UFS resp. model candidates of the FLP reduct) and the number of external atom evaluations.

We evaluated the implementation on a Linux server with two 12-core AMD 6176 SE CPUs with 128GB RAM. In each test run the CPU usage was limited to two CPU cores, running a Condor load distribution system which ensures robust runtimes (i.e., multiple runs of the same instance have negligible deviations). The timeout was uniformly set to 300 seconds for each instance; for each parameter value, the average runtime over all instances is printed where timeouts, whose number is shown in parentheses, are fully taken into account.

6.1 Detailed Benchmark Description

Set Partitioning. This benchmark extends the program from Example 2.5 by the additional constraint $\leftarrow sel(X), sel(Y), sel(Z), X \neq Y, X \neq Z, Y \neq Z$ and varies the size of *domain*. Here we see a big advantage of the UFS check over the explicit check, for both computing all answer sets and finding the first one. A

domain	all answer sets						first answer set					
	explicit	+EBL	UFS Γ mol	UFS Γ mod	+EBL	Ω	explicit	+EBL	UFS Γ mol	UFS Γ mod	+EBL	Ω
5	12.52	6.43	0.73	0.73	0.12	0.11	0.79	0.85	0.08	0.07	0.07	0.10
6	100.48	51.51	1.96	1.51	0.22	0.16	4.72	4.91	0.09	0.07	0.09	0.09
7	—	—	4.80	4.60	0.34	0.20	27.51	29.67	0.13	0.07	0.08	0.10
8	—	—	10.40	11.51	0.49	0.21	173.39	186.59	0.12	0.10	0.12	0.11
9	—	—	28.64	29.23	0.66	0.30	—	—	0.14	0.12	0.16	0.16
10	—	—	65.17	71.17	0.81	0.35	—	—	0.21	0.13	0.20	0.18
11	—	—	161.98	175.03	1.03	0.51	—	—	0.25	0.16	0.23	0.21
12	—	—	—	—	1.48	0.60	—	—	0.27	0.16	0.26	0.25
15	—	—	—	—	3.13	1.04	—	—	0.38	0.26	0.42	0.34
20	—	—	—	—	13.89	2.55	—	—	0.87	0.53	0.59	0.76
25	—	—	—	—	54.81	5.04	—	—	1.84	0.87	1.65	1.50

Table 1: Set Partitioning Benchmark Results

closer investigation shows that the improvement is mainly due to the optimizations in Section 4, which make the UFS check investigate significantly fewer candidates than the explicit FLP check. Furthermore the UFS check requires fewer external computations.

Both the explicit and the UFS-based minimality check benefit from EBL if we compute all answer sets, but the results show that the UFS-based check benefits more. In contrast, UFL (not shown in the table) does not lead to a further speedup as no unfounded sets are found in this program.

The decision criterion and program decomposition improve the runtime slightly for small instances. However, for large instances the decision criterion cannot avoid the UFS check in most components of the program because of its cyclic structure. Thus a single UFS check over the whole program is replaced by multiple UFS checks over individual program components, which involves more overhead that becomes visible when computing all answer sets.

If we compute only one answer set, then EBL turns out to be counter-productive. This is because learning is involved with additional overhead, while we cannot profit much from the learned knowledge if we abort after the first answer set, hence the costs exceed the benefit.

Using the encoding Ω instead of Γ increases the efficiency in this case, because there is not only a large number of answer sets but also a large number of answer set candidates. Thus, a reusable encoding is very beneficial, even if we compute only one answer set.

Multi-Context Systems (MCSs). Nonmonotonic Multi-Context-Systems (MCSs) [Brewka and Eiter, 2007] are a generic formalism for aligning knowledge bases called *contexts*, which emerged from an approach in [Ghidini and Giunchiglia, 2001]. The contexts are interlinked via bridge rules which enable belief exchange across contexts; the MCS semantics requires that local belief sets are compliant with the bridge rules. Such compliance can be impossible to achieve; that is, the MCS is inconsistent. To understand the reasons for the latter, in [Eiter et al., 2012e] *inconsistency explanations (IEs)* for a MCS were defined, which can be computed with a HEX-program encoding. This encoding is based on *Saturation*, which is a general technique for solving Σ_2^p problems in disjunctive answer set programming, cf. [Leone et al., 2006]. Intuitively, a quantified Boolean formula (QBF) of the form $\exists \mathbf{X} \forall \mathbf{Y} \Phi(\mathbf{X}, \mathbf{Y})$ is evaluated using this technique as follows. Disjunctions are used to guess whether a variable v is true or false, where the literals v and $\neg v$ are encoded with designated atoms. A ‘spoil’ atom is made true whenever Φ evaluates to true given truth assignments to \mathbf{X} and \mathbf{Y} . Finally whenever ‘spoil’ is true, all literals over \mathbf{Y} are set to true; this creates a unique assignment of the respective atoms. Given a guess on \mathbf{X} , its unique ‘spoil’ extension is an answer set if and only if

#ctx	explicit	+EBL	UFS Γ mol	UFS Γ mod	+EBL	+UFL	Ω
3 (6)	6.98 (0)	4.45 (0)	4.47 (0)	4.32 (0)	1.62 (0)	0.10 (0)	0.11 (0)
4 (10)	63.75 (1)	47.85 (1)	55.14 (1)	54.90 (1)	26.38 (0)	0.12 (0)	0.13 (0)
5 (8)	177.93 (3)	143.26 (3)	183.79 (4)	185.44 (4)	98.06 (1)	0.13 (0)	0.15 (0)
6 (6)	169.07 (3)	155.16 (3)	178.11 (3)	174.46 (3)	151.87 (3)	0.18 (0)	0.17 (0)
7 (12)	240.13 (9)	227.92 (9)	248.74 (9)	254.15 (10)	215.89 (8)	0.19 (0)	0.22 (0)
8 (5)	250.42 (4)	207.05 (3)	251.19 (4)	251.66 (4)	188.79 (3)	0.22 (0)	0.24 (0)
9 (8)	300.00 (8)	279.44 (7)	300.00 (8)	300.00 (8)	256.77 (6)	0.25 (0)	0.27 (0)
10 (11)	300.00 (11)	271.64 (9)	300.00 (11)	300.00 (11)	244.39 (8)	0.29 (0)	0.32 (0)

Table 2: Consistent MCSs Benchmark Results

#ctx	explicit	+EBL	all answer sets		+EBL	+UFL	Ω
			UFS Γ mol	UFS Γ mod			
3 (9)	4.90 (0)	3.07 (0)	3.46 (0)	3.28 (0)	1.21 (0)	0.17 (0)	0.14 (0)
4 (14)	54.40 (1)	20.35 (0)	46.62 (1)	45.50 (1)	7.06 (0)	3.23 (0)	0.87 (0)
5 (11)	168.68 (5)	149.90 (5)	165.52 (5)	167.53 (5)	104.03 (2)	3.70 (0)	0.88 (0)
6 (18)	138.09 (7)	102.97 (6)	136.29 (7)	136.00 (7)	87.82 (4)	10.52 (0)	2.18 (0)
7 (16)	216.84 (8)	152.41 (8)	201.72 (8)	202.90 (8)	143.93 (7)	120.91 (4)	64.49 (1)
8 (6)	268.62 (5)	149.15 (2)	268.65 (5)	268.48 (5)	116.46 (2)	55.89 (1)	51.38 (1)
9 (14)	291.86 (13)	207.49 (9)	294.29 (13)	294.23 (13)	188.95 (8)	123.70 (5)	110.28 (5)
10 (12)	300.00 (12)	300.00 (12)	300.00 (12)	300.00 (12)	289.56 (11)	290.33 (11)	277.85 (11)

#ctx	explicit	+EBL	first answer set		+EBL	+UFL	Ω
			UFS Γ mol	UFS Γ mod			
3 (9)	0.11 (0)	0.11 (0)	0.10 (0)	0.10 (0)	0.10 (0)	0.10 (0)	0.10 (0)
4 (14)	0.16 (0)	0.16 (0)	0.14 (0)	0.13 (0)	0.13 (0)	0.14 (0)	0.15 (0)
5 (11)	0.18 (0)	0.18 (0)	0.16 (0)	0.15 (0)	0.15 (0)	0.15 (0)	0.16 (0)
6 (18)	0.21 (0)	0.20 (0)	0.18 (0)	0.18 (0)	0.17 (0)	0.17 (0)	0.19 (0)
7 (16)	0.23 (0)	0.21 (0)	0.21 (0)	0.19 (0)	0.16 (0)	0.17 (0)	0.19 (0)
8 (6)	0.26 (0)	0.22 (0)	0.22 (0)	0.22 (0)	0.19 (0)	0.18 (0)	0.20 (0)
9 (14)	0.37 (0)	0.30 (0)	0.33 (0)	0.32 (0)	0.23 (0)	0.24 (0)	0.26 (0)
10 (12)	0.51 (0)	0.36 (0)	0.45 (0)	0.44 (0)	0.28 (0)	0.29 (0)	0.32 (0)

Table 3: Inconsistent MCSs Benchmark Results

all guesses of truth assignments \mathbf{Y} make the spoil atom true and *saturate* the guess to become the unique extension—this holds due to the minimality condition on answer sets of the reduct (see Definition 2.7).

We use the HEX-encoding for computing IEs as a benchmark, as the saturation is rich in cycles through external atoms and disjunctive rule heads. External atoms in this benchmark evaluate semantics of contexts in the MCS (i.e., the local belief sets or models).

We use random instances of different MCS topologies, i.e., connection graphs of contexts, created with our MCS benchmark generator.² Note that the topologies are by their structure bound to certain system sizes (number of contexts), and that the difficulty of the instances varies among topologies; thus larger instances may have shorter runtimes. Our instances have up to 10 contexts, each consisting of a randomly generated consistent normal answer set program.

²Described at <http://www.kr.tuwien.ac.at/research/systems/dmcs/experiments.html>, online available at <https://dmcs.svn.sourceforge.net/svnroot/dmcs/dmcs/trunk>

The number of candidates for smaller models of the FLP reduct equals the number of unfounded set candidates as each unfounded set corresponds to a smaller model. However, as we stop the enumeration as soon as a smaller model respectively an unfounded set is found, the explicit and the UFS check may consider depending on the specific program and solver heuristics different numbers of interpretations. This explains why the UFS check is sometimes slightly slower than the explicit check. However, the delay between different UFS candidates was always smaller, which sometimes makes it faster even if it visits more candidates.

The results for consistent and inconsistent MCSs are shown in Tables 2 and 3, respectively, where the number of instances of of each system size is given in parentheses. Intuitively, consistent and inconsistent MCSs are dual, as for each candidate the explicit resp. UFS check fails (i.e., stops early), vs. for some (or many) candidates the check succeeds (stops late). However, the mixed results do not permit us to draw solid conclusions on the computational relationship of the evaluation methods. Nonetheless, we can see that the UFS check based on Ω was often much faster than the explicit check (up to three orders of magnitude).

As consistent MCSs have no IEs and hence no answer sets, we need not distinguish for them between computing one or all answer sets. The effects of external behavior learning [Eiter et al., 2012a] and of unfounded set learning are clearly evident in the MCS benchmarks, for both computing the first and all answer sets. The UFS check profits more from EBL than the explicit check, further adding to its advantage. By activating UFL (which is not possible in the explicit check) we gain another significant speedup.

We now discuss the effects of the syntactic decision criterion and program decomposition. Due to saturation, the encoding contains cycles where nearly all cycles in the HEX-program contain at least one external atom. Therefore the decision criterion can rarely reduce the set of atoms for which the UFS check needs to be performed. However, the benchmark results for MCS instances confirm that the syntactic check is cheap and does not hurt performance. Over all 186 instances, the total runtime with decision criterion and program decomposition was 11,695 seconds compared to 11,702 seconds without, and the number of instance timeouts was the same.

If we use encoding Ω instead of Γ , we can observe another significant speedup for computing all IEs of inconsistent MCSs. This is because there usually exist many answer sets (often thousands), and thus a reusable encoding is very beneficial. In contrast, if we compute only the first answer set or the MCS is consistent (no answer set exists), then the check with the more involved encoding Ω is slightly slower; its reusability does not pay off if we abort after the first answer set.

In summary, we can observe that the encoding Ω leads to a significant performance gain over encoding Γ , while the decision criterion and decomposition do not help. In our next benchmark we will observe opposite effects.

Abstract Argumentation. In this benchmark we compute ideal set extensions [Dung et al., 2007] for randomly generated instances of abstract argumentation frameworks (AFs) [Dung, 1995] of different sizes. The problem of checking whether a given set of arguments is an ideal set of an AF is co-NP-complete [Dunne, 2009]. In this benchmark we use a HEX encoding that mirrors this complexity: it guesses such a set and checks its ideality using the Saturation technique involving an external atom (see Appendix A.1).

Table 4 shows the results for different numbers of arguments, where each entry is the average of 30 benchmark instances. We compare the following configurations for both computing all and the first answer set.

In the first column we do an explicit minimality check without learning techniques. The second column shows that learning (EBL) leads to almost the same runtime results. This can be explained by the structure of the encoding, which does not allow for effectively reusing learned nogoods.

In the third column, we perform an UFS-based minimality check using our encoding Γ , but without applying the decision criterion and decomposition. We can observe that this is already a significant improvements

#args	explicit	+EBL	all answer sets			Ω
			UFS Γ mol	UFS Γ mod	+EBL +UFL	
5	0.63 (0)	0.64 (0)	0.51 (0)	0.34 (0)	0.34 (0)	0.35 (0)
6	1.35 (0)	1.35 (0)	1.10 (0)	0.66 (0)	0.65 (0)	0.63 (0)
7	3.05 (0)	3.19 (0)	2.38 (0)	0.99 (0)	0.97 (0)	0.98 (0)
8	10.10 (0)	10.15 (0)	5.87 (0)	2.02 (0)	1.94 (0)	1.92 (0)
9	33.76 (0)	34.82 (0)	18.59 (0)	3.45 (0)	3.35 (0)	3.35 (0)
10	109.58 (1)	112.57 (1)	49.84 (0)	8.86 (0)	8.32 (0)	7.92 (0)
15	290.01 (29)	290.01 (29)	290.00 (29)	108.64 (6)	99.88 (6)	98.44 (6)
20	300.00 (30)	300.00 (30)	300.00 (30)	185.94 (15)	180.25 (15)	179.60 (14)

#args	explicit	+EBL	first answer set			Ω
			UFS Γ mol	UFS Γ mod	+EBL +UFL	
5	0.41 (0)	0.41 (0)	0.41 (0)	0.27 (0)	0.27 (0)	0.28 (0)
6	0.87 (0)	0.87 (0)	0.82 (0)	0.54 (0)	0.53 (0)	0.54 (0)
7	1.44 (0)	1.48 (0)	1.27 (0)	0.66 (0)	0.66 (0)	0.67 (0)
8	4.74 (0)	4.86 (0)	3.59 (0)	1.49 (0)	1.48 (0)	1.44 (0)
9	17.47 (0)	17.74 (0)	11.88 (0)	2.23 (0)	2.22 (0)	2.18 (0)
10	52.22 (0)	54.18 (0)	29.17 (0)	4.08 (0)	3.88 (0)	3.62 (0)
15	260.03 (26)	260.03 (26)	260.02 (26)	33.13 (2)	32.85 (2)	32.36 (2)
20	220.11 (22)	220.11 (22)	220.04 (22)	72.30 (4)	64.28 (3)	64.50 (3)

Table 4: Argumentation Benchmark Results

compared to the explicit minimality check, illustrating the effectiveness of our new approach. Similar as in the MCS benchmark, the number of reduct model candidates is equal to the number of UFS candidates in most cases, but the UFS check again enumerates its candidates faster; this explains the observed speedup.

When we enable the decision criterion and program decomposition, we can observe a further speedup. This is because cycles in argumentation instances usually involve only small parts of the overall program; thus the UFS search can be significantly simplified by excluding large program parts. We further have observed that program decomposition without the decision criterion is counter-productive (not shown in the table), because a single UFS search over the whole program is replaced by many UFS searches over program components (without the decision criterion, no such check is excluded). This incurs more overhead.

In the fifth column we enable EBL and UFL, which leads to a small speedup in some cases. However, as already mentioned above, no effective reuse of learned nogoods is possible.

Finally, switching the encoding from Γ to Ω leads to a small speedup in some cases, but is also counter-productive in others. This is because the programs in this benchmark have usually only very few compatible sets, and only few unfounded set checks need to be performed. Hence the lower initialization overhead of the encoding Ω does not influence the runtime dramatically. On the other hand, the higher complexity of the encoding Ω increases the runtime of small instances.

Default Reasoning over Description Logics Benchmarks.. A prominent instance of HEX-programs are DL-programs, which combine description logic ontologies with rules; they result by using a special external atom that is available as DL-plugin in dlvhex. Via DL-programs, we obtain an encoding of terminological default reasoning over description logic ontologies in the approach of [Baader and Hollunder, 1995] into HEX-programs, in which defaults require cyclic dependencies over external atoms. However, as all such dependencies involve default negated atoms, we have no cycles according to Definition 5.1, which respects only positive dependencies. Hence, the decision criterion comes to the conclusion that no UFS check is

required.

We used variants of the benchmarks presented in [Eiter et al., 2012a], which query wines from an ontology and classify them as red or white wines, where a wine is assumed to be white unless the ontology explicitly entails the opposite. In this scenario, the decision criterion eliminates all unfounded set checks. However, as there is only one compatible set per instance, there would be only one unfounded set check anyway, hence the speedup due to the decision criterion is not significant. But the effect of the decision criterion can be increased by slightly modifying the scenario such that there are multiple compatible sets. This can be done, for instance, by nondeterministic default classifications, e.g., if a wine is not Italian, then it is either French or Spanish by default. Our experiments have shown that with a small number of compatible sets, the performance enhancement due to the decision criterion is marginal, but increases with larger numbers of compatible sets. For instance, for 243 compatible sets (and thus 243 unfounded set checks) we could observe a speedup from 13.59 to 12.19 seconds.

Conformant Planning. In classical AI planning, a planning domain contains a description of actions with their preconditions and effects in the world. Finding a plan means to find a sequence of actions that reaches from a given initial state a state fulfilling a given goal condition. Conformant planning [Goldman and Boddy, 1996, Smith and Weld, 1998] is the same problem but where the initial state is only partially specified and/or the domain is nondeterministic, such that by executing the plan we reach the goal regardless of the action outcomes and the actual initial state.

We here experimented on a very simple conformant planning domain: two robots with a limited sensor range patrol an area, in which is an object at an unknown initial location. The goal is to find a sequence of movements of the two robots such that they detect the object in all cases. For experiments we used an encoding which realizes conformant planning using Saturation (see above) and contains an external atom for computing whether the patrol robots detect object (cf. Appendix A.2. In general, deciding the existence of a short (polynomial length bounded) conformant plan is Σ_3^P -complete, see [Turner, 2002], but if action executability is decidable in polynomial time, the problem is in Σ_2^P ; our example domain enjoys the property.

Our results are displayed in Table 5, which shows averages over 10 instances per size. The instances consist of $n \times 4$ grids with $n \in \{3, \dots, 16\}$, the plan length required for finding a solution increases with larger instance sizes. (This is because the number of robots does not increase while the two robots must still cover the whole area.) Instances were generated by randomly placing robots in opposite quarters of the map.

As expected we observe that the explicit FLP check performs worst, followed by the monolithic UFS check with Γ encoding, and the modular UFS Γ encoding; the UFS Ω encoding (without external behavior nor unfounded set learning) performs best. External behavior learning (EBL) and unfounded set learning (UFL) do not improve the performance, on the contrary it increases the run times significantly for the modular UFS Γ check and slightly for the Ω check. EBL does not change times significantly for the explicit check, therefore we omit results for explicit +EBL.

By looking into profiling information and at the domain we found the reasons: (a) the external atoms depend on a large part of the interpretation (locations of all robots) so EBL cannot cut away significant portions of the search space; (b) evaluating the external atom takes a negligible amount of time, so beneficial effects of EBL will be outweighed by its computational overhead. For UFL we observed that the benchmark instances contain only few unfounded sets (pruning less than half of the answer set candidates) and thus UFL cannot improve performance given the overhead it incurs. We conclude that in some scenarios, using EBL and UFL can reduce efficiency.

As in the Ω check the UFS check encoding is constructed only once, the overhead for EBL and UFL observed with the Γ encoding no longer has such a big impact. Nevertheless Ω without learning slightly

map size	plan length	all answer sets						
		explicit	UFS Γ mol	UFS Γ mod	+EBL	+UFL	Ω	+EBL+UFL
3×4	1	12.67 (0)	0.90 (0)	0.66 (0)	0.88 (0)	0.84 (0)	0.75 (0)	0.75 (0)
4×4	1	21.31 (0)	0.85 (0)	0.74 (0)	0.91 (0)	0.63 (0)	0.77 (0)	0.70 (0)
5×4	1	27.33 (0)	0.84 (0)	0.73 (0)	0.82 (0)	0.82 (0)	0.68 (0)	0.55 (0)
6×4	2	272.39 (7)	4.64 (0)	2.62 (0)	3.36 (0)	3.51 (0)	1.92 (0)	2.02 (0)
7×4	2	291.26 (9)	6.80 (0)	3.30 (0)	5.14 (0)	5.37 (0)	2.67 (0)	3.23 (0)
8×4	3	300.00 (10)	107.87 (0)	33.73 (0)	78.39 (0)	79.08 (0)	16.59 (0)	22.03 (0)
9×4	3	300.00 (10)	124.72 (0)	40.61 (0)	104.94 (0)	107.37 (0)	21.14 (0)	28.61 (0)
10×4	4	300.00 (10)	300.00 (10)	300.00 (10)	300.00 (10)	300.00 (10)	230.87 (2)	283.29 (8)
11×4	4	300.00 (10)	300.00 (10)	295.51 (9)	300.00 (10)	300.00 (10)	209.69 (3)	275.22 (6)
12×4	5	300.00 (10)	300.00 (10)	300.00 (10)	300.00 (10)	300.00 (10)	294.98 (9)	300.00 (10)
13×4	5	300.00 (10)	300.00 (10)	300.00 (10)	300.00 (10)	300.00 (10)	295.37 (8)	300.00 (10)
14×4	6	300.00 (10)	300.00 (10)	300.00 (10)	300.00 (10)	300.00 (10)	293.65 (9)	300.00 (10)
15×4	6	300.00 (10)	300.00 (10)	300.00 (10)	300.00 (10)	300.00 (10)	300.00 (10)	300.00 (10)
16×4	7	300.00 (10)	300.00 (10)	300.00 (10)	300.00 (10)	300.00 (10)	300.00 (10)	300.00 (10)

map size	plan length	first answer set						
		explicit	UFS Γ mol	UFS Γ mod	+EBL	+UFL	Ω	+EBL+UFL
3×4	1	1.44 (0)	0.05 (0)	0.04 (0)	0.04 (0)	0.04 (0)	0.05 (0)	0.05 (0)
4×4	1	2.21 (0)	0.06 (0)	0.04 (0)	0.05 (0)	0.04 (0)	0.05 (0)	0.06 (0)
5×4	1	6.06 (0)	0.08 (0)	0.06 (0)	0.07 (0)	0.06 (0)	0.07 (0)	0.07 (0)
6×4	2	15.46 (0)	0.23 (0)	0.13 (0)	0.13 (0)	0.13 (0)	0.12 (0)	0.14 (0)
7×4	2	35.43 (0)	0.43 (0)	0.20 (0)	0.24 (0)	0.24 (0)	0.19 (0)	0.22 (0)
8×4	3	188.35 (3)	7.46 (0)	2.13 (0)	1.93 (0)	1.93 (0)	1.25 (0)	1.74 (0)
9×4	3	181.38 (3)	9.98 (0)	3.18 (0)	2.50 (0)	3.02 (0)	1.48 (0)	1.84 (0)
10×4	4	248.52 (8)	83.65 (2)	51.98 (1)	55.91 (1)	62.08 (1)	36.50 (1)	40.44 (1)
11×4	4	300.00 (10)	117.33 (2)	51.74 (1)	67.10 (1)	84.32 (2)	37.00 (1)	43.48 (1)
12×4	5	274.38 (9)	254.15 (7)	176.32 (4)	216.24 (6)	198.95 (5)	94.64 (2)	125.79 (2)
13×4	5	300.00 (10)	293.08 (9)	240.73 (5)	276.55 (9)	276.53 (9)	164.91 (4)	203.02 (4)
14×4	6	300.00 (10)	286.46 (9)	247.39 (8)	257.14 (8)	257.27 (8)	207.95 (6)	205.80 (6)
15×4	6	300.00 (10)	300.00 (10)	281.90 (9)	300.00 (10)	300.00 (10)	252.76 (8)	277.33 (9)
16×4	7	300.00 (10)	300.00 (10)	300.00 (10)	300.00 (10)	300.00 (10)	292.24 (9)	300.00 (10)

Table 5: Conformant Planning Benchmark Results

outperforms Ω with UFL and EBL. An analysis of the number of UFS checks and the number of external atom evaluations (not shown in table) revealed that UFL and EBL decreases (i) the number of unfounded set checks needed and (ii) the number of external atoms evaluated (in one 9×4 instance, from 5132 to 3341). Thus if external computations would be more costly, the positive effects of UFL and EBL on Ω would outweigh their computational overhead and it would be beneficial to activate UFL and EBL in this domain.

Interestingly, for small map sizes we see that with the Ω encoding, the 3×4 instances actually seem to be harder to solve than the larger 4×4 and 5×4 instances. This is because all these instances require plan length 1 to find a solution; so the larger instances are more constrained than the smaller instances because there robots have less freedom to move around while still detecting all objects. Hence, for 5×4 maps the solver finds solutions faster than for 4×4 maps.

The conclusion from this benchmark is that depending on the computational task in external atoms, UFL and EBL can be beneficial or harmful for efficiency of reasoning.

6.2 Summary

Our experiments have shown that the learning technique EBL developed in [Eiter et al., 2012a] and the techniques introduced in this paper lead to significant performance improvements in most cases, with few exceptions for specific benchmarks. The effects of external behavior learning (EBL) are clearly evident both for the explicit minimality check and for the unfounded set-based check, but are even more prominent with the latter. Independently of whether EBL is used or not, unfounded set checking pushes the efficiency of HEX-program evaluation compared to explicit minimality checking. Moreover, it allows for learning of additional nogoods, which is also advantageous in most of our benchmarks. Regarding the two problem encodings, the benchmarks show that the UFS check is usually faster with the Ω encoding than with the Γ encoding, however the former UFS check involves more initialization overhead, which might be counterproductive for small programs.

The decision criterion may lead to an additional speedup and does not introduce notable overhead, thus it can always be activated. Finally, program decomposition often leads to an additional performance gain, but should only be used in combination with the decision criterion because otherwise a single UFS check is replaced by multiple UFS checks, which involves more overhead.

7 Discussion

7.1 Related Work

Foundedness is besides grounding one of the main differences between ASP and SAT solving. The notion of unfounded set has been introduced by van Gelder et al. 1991 for defining the well-founded semantics of logic programs with negation, by constructing the least fixpoint of a monotone operator on partial interpretations; the total fixpoints of this operator correspond to the stable models of the logic program. This actually allows to give a characterization of the stable models in terms of unfounded sets. In fact, unfounded set checking turned out to be a fruitful model-based approach in ASP solving, which has an equally successful syntactic counterpart known as *loop formulas* [Lin and Zhao, 2004, Lee and Lifschitz, 2003, Lee, 2005]. Different kinds of unfounded set checks with different complexities have been developed for various program classes.

The computation of answer sets by a model generate and test approach, which is pursued by many ASP solvers, requires that a form of minimality check or unfounded set check is carried out already for ordinary logic programs (i.e., in absence of external atoms). However, for normal program this test is tractable and it is frequently realized using *source pointers* [Simons et al., 2002]. Intuitively, the reasoner stores for each atom a pointer to a rule which possibly supports this atom. The list of source pointers is updated during propagation. If at some point there is no supporting rule for an atom, then it can conclude that this atom must be false.

In the context of ASP, the notion of unfounded set has been explicitly formulated and extended to disjunctive logic programs by Leone et al. 1997, who proved that the stable models of a disjunctive logic program are its models that are unfounded-free. This results was the basis for the architecture of the dlv solver, which generates answer set candidates that are checked for unfounded-freeness. This test, which like answer set checking from FLP-reducts is co-NP-complete [Faber et al., 2011], has been reduced in [Koch et al., 2003] to unsatisfiability testing of a SAT instance. This approach has been later extended to conflict-driven learning and unfounded set checking in [Drescher et al., 2008], where two instances for the CLASP solver, an ASP instance and a SAT instance, are used to generate and check answer set candidates, respectively. Alviano et al. 2011 consider normal logic programs with monotone and antimonotone aggregate atoms, and defined unfounded sets for such programs. Based on this, they extended the well-founded semantics of Van

Gelder et al. 1991, which—although closely related—is weaker than the general FLP semantics [Faber et al., 2011].

We now considered unfounded set checking in the presence of external sources and for FLP answer sets, for which the results in [Drescher et al., 2008] do not immediately carry over and also the problem is intuitively computationally harder. Indeed, already for Horn programs, the presence of nonmonotonic external atoms that are decidable in polynomial time makes unfoundedness checking intractable (more precisely, co-NP-complete), such that deciding the existence of an FLP answer set is a Σ_2^P -complete problem in the ground case; for computationally harder external atoms, the complexity may increase relative to an oracle for the external function (see [Faber et al., 2011]).

Drescher et al. 2008 employed also a splitting technique related to our program decomposition, but for ordinary programs only. Note that our notion of splitting is different from the well-known splitting sets in [Lifschitz and Turner, 1994], as we consider only positive dependencies for ordinary atoms. While we consider e-cycles, which are specific for HEX-programs, the interest in [Drescher et al., 2008] was with head-cycles, which may arise with disjunctive rule heads. In fact, our approach may be regarded as an extension of the one of Drescher et al., since the evaluation of $\hat{\Pi}$ follows their principles of performing UFS checks in case of head-cycles.

The unfounded set check that we consider in this work is instrumental for the FLP semantics (as it is for the standard GL-reduct based semantics); we just mention that other semantics of HEX-programs may not involve such a step. For instance, [Shen, 2011] and [Shen and Wang, 2011] present a well-justified semantics where unfounded set checking is essentially replaced by a fixpoint iteration which, intuitively, tests if a model candidate reproduces itself but excludes circular justifications. However, the complexity of answer set computation does not decrease by this approach in general, and in particular the Σ_2^P -completeness of deciding answer set existence of ground Horn programs with nonmonotonic external atoms that are decidable in polynomial time holds.

7.2 Extensions

We have designed our unfounded set check as a postcondition test; that is like the explicit FLP check, it is carried out only after a complete interpretation has been generated as an answer set candidate. However, in certain cases it might be already clear that a partial interpretation (in which some truth values are open) can never become an answer set, because the existence of an unfounded set is guaranteed for any completion to a total interpretation. One can then backtrack earlier, which intuitively leads to a saving for certain classes of instances.

Exploring this idea, we have generalized our framework with a control component which decides, based on a heuristics, when an unfounded set check is carried out; in the standard setting, this is whenever an assignment in the model generation is complete (i.e., a total interpretation is given).

A UFS check on partial interpretations that is sound with respect to any completion to a total interpretation is possible if the ASP solver has finished internal unit propagation over a maximal subset of the program such that the interpretation is already complete on it. We thus used this criterion, which is easy to test, for a greedy heuristics to issue UFS checks in our prototype system.

However, in contrast to our initial expectation, we found that for all our benchmarks the UFS check over partial interpretations was not productive. A closer look reveals that this is essentially because nogood learning from unfounded sets (UFL) effectively avoids the reconstruction of the same unfounded set anyway. It therefore rarely happens that UFS checking over a partial interpretation identifies an unfounded set earlier than UFS checking over complete interpretations. As UFS checking for HEX-programs involves the evaluation

of external sources and compatibility testing, this easily leads to costs that are higher than the potential savings. On the other hand, for ordinary ASP solving (where no such extra costs incur), UFS checks over partial interpretations may still be beneficial. In conclusion, UFS checks on partial interpretations of HEX-programs will require tailored heuristics that not only take the structure of the program, but also domain-specific knowledge into account, which remains for future work.

8 Conclusion

HEX-programs are an expressive extension of non-monotonic logic programs with access to external information via external atoms, which supported by a plugin architecture can be fruitfully deployed for a range of applications. External atoms however make the efficient evaluation of HEX-programs a challenging task, and in particular to compute answer sets of a HEX-program Π , which are the models \mathbf{A} of Π that are subset-minimal models of its FLP-reduct $f\Pi^{\mathbf{A}}$ (which keeps all rules whose bodies are satisfied). To improve on this expensive test (which was customary in implementations so far), we have presented an alternative test based on unfounded sets that we obtain by adapting the notion of unfounded set for aggregates in [Faber, 2005] to external atoms. Also [Alviano et al., 2011] use a related notion of unfounded sets for programs with aggregates, but restrict the discussion to monotonic and antimonotonic aggregates. We have realized unfounded set (UFS) checking by a transformation to SAT solving, where the satisfying assignments of a constructed CNF generate candidate unfounded sets, which in turn are subject to a (rather simple) postcheck that takes external atom evaluation into account.

In particular, we have provided two SAT encodings for UFS checking: the conceptually simple encoding $\Gamma_{\Pi}^{\mathbf{A}}$, which needs initialization for every UFS check, and the advanced encoding Ω_{Π} , which can be reused for all UFS checks. To further boost performance, we have shown how to learn from unfounded sets by deriving nogoods, i.e., constraints (possibly involving also external atoms) which guide future search in model generation and help to avoid that unfounded sets are regenerated. In further elaboration, we have refined the basic approach by suitable program splitting, such that the UFS check can be carried out independently on program components, cutting down the complexity. Furthermore, we have presented a syntactic criterion that allows us to decide efficiently whether the UFS check can be safely skipped for a component or the whole program, exploiting that the answer set candidates from the model search have only special unfounded sets that involve cyclic input to external atoms; for HEX-programs in simple applications, this is usually not the case.

The experimental evaluation of our new approach, which considered different combinations of the techniques and comprised problems from various domains including multi-context systems, abstract argumentation, default reasoning over ontologies, and conformant planning, where HEX-programs serve for easy-cut declarative problem solving, has shown that it is much more efficient than the traditional minimal model check; it can lead to exponential gains and yields often drastic performance improvements, while it is not slower (except in very few cases by a marginal amount). Furthermore, the reusable encoding Ω_{Π} turned out to be beneficial for programs that require many unfounded set checks, which includes all programs that have many answer sets.

Future work. An issue for further improvement of the approach in this paper are other heuristics for UFS checking over partial interpretations. While the natural heuristics that we considered are counterproductive, others might lead to additional speedup in some cases. However, this may require to incorporate domain-specific knowledge about external atoms; currently this is only done in the learning algorithms but not in the heuristics. In the same line, one might consider developing a heuristics for dynamically choosing

between the UFS search encodings that we presented, and to study heuristics for guiding the unfounded set search, i.e., for variable selection by the SAT solver. Currently, our implementation applies the same heuristics for the unfounded set search as in the model generation process. However, our experimental comparison with the explicit minimality check in terms of the considered candidate sets suggests that there is room for improvement by employing suitable choices. Developing appropriate heuristics and validating their effectiveness on candidate set enumeration remains to be explored. Finally, an obvious issue are other criteria that allow to skip the UFS check or simplify it while they can be efficiently tested.

References

- M. Alviano, F. Calimeri, W. Faber, N. Leone, and S. Perri. Unfounded Sets and Well-Founded Semantics of Answer Set Programs with Aggregates. *Journal of Artificial Intelligence Research*, 42:487–527, 2011. doi: 10.1613/jair.3432. URL <http://www.jair.org/papers/paper3432.html>.
- F. Baader and B. Hollunder. Embedding Defaults into Terminological Knowledge Representation Formalisms. *Journal of Automated Reasoning*, 14(1):149–180, 1995. doi: 10.1007/BF00883932.
- S. Basol, O. Erdem, M. Fink, and G. Ianni. HEX Programs with Action Atoms. In M. Hermenegildo and T. Schaub, editors, *Technical Communications of the 26th International Conference on Logic Programming (ICLP'10)*, volume 7 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 24–33, Dagstuhl, Germany, 2010. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. ISBN 978-3-939897-17-0. doi: 10.4230/LIPIcs.ICLP.2010.24. URL <http://drops.dagstuhl.de/opus/volltexte/2010/2580>.
- G. Brewka and T. Eiter. Equilibria in Heterogeneous Nonmonotonic Multi-Context Systems. In R. C. Holte and A. Howe, editors, *22nd AAAI Conference on Artificial Intelligence (AAAI'07)*, pages 385–390. AAAI Press, 2007. URL <http://www.informatik.uni-leipzig.de/~brewka/papers/Equilibria.pdf>.
- G. Brewka, T. Eiter, and M. Truszczyński. Answer set programming at a glance. *Communications of the ACM*, 54(12):92–103, 2011. doi: 10.1145/2043174.2043195.
- C. Drescher, M. Gebser, T. Grote, B. Kaufmann, A. König, M. Ostrowski, and T. Schaub. Conflict-driven disjunctive answer set solving. In G. Brewka and J. Lang, editors, *11th International Conference Principles of Knowledge Representation and Reasoning (KR 2008)*, Sydney, Australia, September 16-19, 2008, pages 422–432. AAAI Press, 2008.
- P. Dung, P. Mancarella, and F. Toni. Computing ideal sceptical argumentation. *Artificial Intelligence*, 171: 642–674, 2007.
- P. M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77(2):321–357, 1995. doi: 10.1016/0004-3702(94)00041-X.
- P. E. Dunne. The computational complexity of ideal semantics. *Artificial Intelligence*, 173(18):1559–1591, 2009.

- U. Egly, S. A. Gaggl, and S. Woltran. Answer-set programming encodings for argumentation frameworks. *Argument and Computation*, 1(2):147–177, 2010.
- T. Eiter, G. Ianni, R. Schindlauer, and H. Tompits. A Uniform Integration of Higher-Order Reasoning and External Evaluations in Answer-Set Programming. In L. P. Kaelbling and A. Saffiotti, editors, *19th International Joint Conference on Artificial Intelligence (IJCAI'05)*, pages 90–96. Professional Book Center, 2005. ISBN 0938075934. URL http://www.kr.tuwien.ac.at/staff/eiter/et-archive/ijcai05-hex_asp.pdf.
- T. Eiter, G. Ianni, R. Schindlauer, and H. Tompits. Effective Integration of Declarative Rules with External Evaluations for Semantic-Web Reasoning. In Y. Sure and J. Domingue, editors, *3rd European Conference on Semantic Web (ESWC'06)*, volume 4011 of *LNCS*, pages 273–287. Springer, 2006. doi: 10.1007/11762256_22.
- T. Eiter, G. Ianni, T. Krennwallner, and R. Schindlauer. Exploiting conjunctive queries in description logic programs. *Annals of Mathematics and Artificial Intelligence*, 53(1–4):115–152, August 2008a. ISSN 1012-2443. doi: 10.1007/s10472-009-9111-3. URL <http://www.kr.tuwien.ac.at/staff/tkren/pub/2008/amai2008-cqprog.pdf>.
- T. Eiter, G. Ianni, T. Lukasiewicz, R. Schindlauer, and H. Tompits. Combining answer set programming with description logics for the semantic web. *Artificial Intelligence*, 172(12-13):1495–1539, 2008b. doi: 10.1016/j.artint.2008.04.002.
- T. Eiter, M. Fink, T. Krennwallner, and C. Redl. Conflict-driven ASP Solving with External Sources. *Theory and Practice of Logic Programming*, 12(4-5):659–679, July 2012a.
- T. Eiter, M. Fink, T. Krennwallner, C. Redl, and P. Schüller. Eliminating Unfounded Set Checking for HEX-Programs. In M. Fink and Y. Lierler, editors, *5th Workshop on Answer Set Programming and Other Computing Paradigms (ASPOCP 2012), September 4, 2012, Budapest, Hungary*, pages 83–97, September 2012b. URL <http://www.kr.tuwien.ac.at/staff/tkren/pub/2012/aspocp2012-elim.pdf>.
- T. Eiter, M. Fink, T. Krennwallner, C. Redl, and P. Schüller. Exploiting Unfounded Sets for HEX-Program Evaluation. In L. F. del Cerro, A. Herzig, and J. Mengin, editors, *13th European Conference on Logics in Artificial Intelligence (JELIA 2012), September 26-28, 2012, Toulouse, France*, volume 7519 of *LNCS*, pages 160–175. Springer, September 2012c. ISBN 978-3-642-33352-1. doi: 10.1007/978-3-642-33353-8_13. URL <http://www.kr.tuwien.ac.at/staff/tkren/pub/2012/jelia2012-ufsflp.pdf>.
- T. Eiter, M. Fink, T. Krennwallner, C. Redl, and P. Schüller. Improving HEX-Program Evaluation based on Unfounded Sets. Technical Report INFSYS RR-1843-12-08, Institut für Informationssysteme, Technische Universität Wien, A-1040 Vienna, Austria, July 2012d.
- T. Eiter, M. Fink, P. Schüller, and A. Weinzierl. Finding explanations of inconsistency in nonmonotonic multi-context systems. Technical Report INFSYS RR-1843-12-09, INFSYS RR-1843-03-08, Inst. für Informationssysteme, TU Wien, December 2012e. Preliminary version in Proc. 12th International Conference on Knowledge Representation and Reasoning (KR 2010), pp. 329–339, AAAI Press, 2010.

- W. Faber. Unfounded sets for disjunctive logic programs with arbitrary aggregates. In C. Baral, G. Greco, N. Leone, and G. Terracina, editors, *8th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR'05)*, volume 3662, pages 40–52. Springer, 2005. doi: 10.1007/11546207_4.
- W. Faber, N. Leone, and G. Pfeifer. Semantics and complexity of recursive aggregates in answer set programming. *Artificial Intelligence*, 175(1):278–298, January 2011. doi: 10.1016/j.artint.2010.04.002.
- M. Gebser, B. Kaufmann, and T. Schaub. Conflict-driven answer set solving: From theory to practice. *Artificial Intelligence*, 187–188:52–89, August 2012. doi: 10.1016/j.artint.2012.04.001.
- M. Gelfond and V. Lifschitz. Classical Negation in Logic Programs and Disjunctive Databases. *New Generation Computing*, 9(3–4):365–386, 1991.
- C. Ghidini and F. Giunchiglia. Local models semantics, or contextual reasoning=locality+compatibility. *Artificial Intelligence*, 127(2):221–259, April 2001. doi: 10.1016/S0004-3702(01)00064-9.
- R. Goldman and M. Boddy. Expressive Planning and Explicit Knowledge. In B. Drabble, editor, *3rd International Conference on Artificial Intelligence Planning Systems (AIPS'96)*, pages 110–117. AAAI Press, 1996. URL <https://www.aaai.org/Papers/AIPS/1996/AIPS96-014.pdf>.
- R. Hoehndorf, F. Loebe, J. Kelso, and H. Herre. Representing default knowledge in biomedical ontologies: application to the integration of anatomy and phenotype ontologies. *BMC Bioinformatics*, 8:377, 2007. doi: 10.1186/1471-2105-8-377.
- C. Koch, N. Leone, and G. Pfeifer. Enhancing disjunctive logic programming systems by SAT checkers. *Artificial Intelligence*, 151(1–2):177–212, 2003. doi: 10.1016/S0004-3702(03)00078-X.
- J. Lee. A model-theoretic counterpart of loop formulas. In L. P. Kaelbling and A. Saffiotti, editors, *19th International Joint Conference on Artificial Intelligence (IJCAI'05)*, pages 503–508. Professional Book Center, 2005. ISBN 0938075934. URL <http://www.ijcai.org/papers/1280.pdf>.
- J. Lee and V. Lifschitz. Loop Formulas for Disjunctive Logic Programs. In C. Palamidessi, editor, *19th International Conference on Logic Programming (ICLP'03)*, volume 2916 of *LNCS*, pages 451–465. Springer, December 2003.
- N. Leone, P. Rullo, and F. Scarcello. Disjunctive Stable Models: Unfounded Sets, Fixpoint Semantics, and Computation. *Information and Computation*, 135(2):69–112, 1997. URL <http://dx.doi.org/10.1006/inco.1997.2630>.
- N. Leone, G. Pfeifer, W. Faber, T. Eiter, G. Gottlob, S. Perri, and F. Scarcello. The DLV System for Knowledge Representation and Reasoning. *ACM Transactions on Computational Logic*, 7(3):499–562, July 2006.
- V. Lifschitz and H. Turner. Splitting a logic program. In P. V. Hentenryck, editor, *11th International Conference on Logic Programming (ICLP'94)*, pages 23–37. MIT Press, 1994. URL <http://www.cs.utexas.edu/users/ai-lab/pub-view.php?PubID=26815>.
- F. Lin and Y. Zhao. ASSAT: computing answer sets of a logic program by SAT solvers. *Artificial Intelligence*, 157(1–2):115–137, 2004. doi: 10.1016/j.artint.2004.04.004.

- M. Marano, P. Obermeier, and A. Polleres. Processing RIF and OWL2RL within DLVHEX. In P. Hitzler and T. Lukasiewicz, editors, *4th International Conference on Web Reasoning and Rule Systems (RR'10)*, volume 6333 of *LNCS*, pages 244–250. Springer, 2010. ISBN 978-3-642-15917-6. doi: 10.1007/978-3-642-15918-3_24.
- D. V. Nieuwenborgh, M. D. Cock, and D. Vermeir. Computing fuzzy answer sets using dlvhex. In V. Dahl and I. Niemelä, editors, *23rd International Conference on Logic Programming (ICLP'07)*, volume 4670 of *LNCS*, pages 449–450. Springer, 2007a. ISBN 978-3-540-74608-9. doi: 10.1007/978-3-540-74610-2_40.
- D. V. Nieuwenborgh, T. Eiter, and D. Vermeir. Conditional planning with external functions. In C. Baral, G. Brewka, and J. S. Schlipf, editors, *9th International Conference on Logic Programming and Non-monotonic Reasoning (LPNMR'07)*, volume 4483 of *LNCS*, pages 214–227. Springer, 2007b. ISBN 978-3-540-72199-4. doi: 10.1007/978-3-540-72200-7_19.
- Y.-D. Shen. Well-supported semantics for description logic programs. In T. Walsh, editor, *22nd International Joint Conference on Artificial Intelligence (IJCAI'11)*, pages 1081–1086. AAAI Press, 2011. doi: 10.5591/978-1-57735-516-8/IJCAI11-185.
- Y.-D. Shen and K. Wang. Extending logic programs with description logic expressions for the semantic web. In L. Aroyo, C. Welty, H. Alani, J. Taylor, A. Bernstein, L. Kagal, N. F. Noy, and E. Blomqvist, editors, *10th International Semantic Web Conference (ISWC'11)*, volume 7031 of *LNCS*, pages 633–648. Springer, 2011. doi: 10.1007/978-3-642-25073-6_40.
- P. Simons, I. Niemelä, and T. Soinen. Extending and implementing the stable model semantics. *Artificial Intelligence*, 138(1-2):181–234, 2002.
- D. E. Smith and D. S. Weld. Conformant Graphplan. In J. Mostow, C. Rich, and B. Buchanan, editors, *15th National Conference on Artificial Intelligence (AAAI'98)*, pages 889–896. AAAI Press / The MIT Press, July 1998.
- H. Turner. Polynomial-length planning spans the polynomial hierarchy. In S. Flesca, S. Greco, N. Leone, and G. Ianni, editors, *European Conference on Logics in Artificial Intelligence (JELIA'02)*, volume 2424 of *LNCS*, pages 111–124. Springer, 2002.
- A. Van Gelder, K. A. Ross, and J. S. Schlipf. The Well-Founded Semantics for General Logic Programs. *Journal of the ACM*, 38(3):619–649, July 1991. doi: 10.1145/116825.116838.
- J. Zakraoui and W. L. Zagler. A logical approach to web user interface adaptation. In A. Holzinger and K.-M. Simoncic, editors, *7th Conference of the Workgroup Human-Computer Interaction and Usability Engineering of the Austrian Computer Society (USAB'11)*, volume 7058 of *LNCS*, pages 645–656. Springer, 2011. ISBN 978-3-642-25363-8. doi: 10.1007/978-3-642-25364-5_46.
- H. Zirtiloğlu and P. Yolum. Ranking semantic information for e-government: complaints management. In *1st International Workshop on Ontology-supported Business Intelligence (OBI'08)*, number 5 in OBI'08, page 7. ACM, 2008. ISBN 978-1-60558-219-1. doi: 10.1145/1452567.1452572.

A Benchmark Encodings

In this appendix, we give details to some of the benchmark encodings (those which are not described in the references). We note that these encodings have not been developed and tuned for good performance and serve merely for an experimental comparison of the various FLP check realizations.

A.1 Abstract Argumentation

The *Abstract Argumentation* benchmark results in Section 6.1 were obtained using the following encoding, which is derived from encodings for admissible and preferred set extensions of an argumentation framework (A, att) described in [Egly et al., 2010].

Input instances of this benchmark are defined over a set A of arguments encoded as facts $arg(a)$ for each $a \in A$ and a set att of attacks between arguments, encoded as facts $att(a, b)$ for some $(a, b) \in A \times A$. The encoding consists of the following rules where $x, y, z \in A$; very similar encodings are explained in detail in [Egly et al., 2010] (but without the use of external atoms).

We define defeat from attacks.

$$defeat(x, y) \leftarrow att(x, y).$$

We guess a set $S \subseteq A$ using predicates in_S and out_S .

$$in_S(x) \leftarrow not\ out_S(x), arg(x). \quad out_S(x) \leftarrow not\ in_S(x), arg(x).$$

We require that all arguments in S are conflict-free and defended from S .

$$\begin{aligned} &\leftarrow in_S(x), in_S(y), defeat(x, y). \\ defeated(x) &\leftarrow in_S(y), defeat(y, x). \\ notDefended(x) &\leftarrow defeat(y, x), not\ defeated(y). \\ &\leftarrow in_S(x), notDefended(x). \end{aligned}$$

For saturation we define a linear order on arguments, including infimum and supremum.

$$\begin{aligned} lt(x, y) &\leftarrow arg(x), arg(y). & (x < y) \\ nsucc(x, z) &\leftarrow lt(x, y), lt(y, z). \\ succ(x, y) &\leftarrow lt(x, y), not\ nsucc(x, y). \\ ninf(x) &\leftarrow lt(y, x). & nsup(x) \leftarrow lt(x, y). \\ inf(x) &\leftarrow not\ ninf(x), arg(x). & sup(x) \leftarrow not\ nsup(x), arg(x). \end{aligned}$$

We perform a guess over a set $T \subseteq A$ using a disjunction.

$$in_T(x) \vee out_T(x) \leftarrow arg(x).$$

We check each argument of T whether it is in S and spoil the answer set if $S \subseteq T$.

$$\begin{aligned}
sInT_{upto}(y) &\leftarrow inf(y), in_S(y), in_T(y). \\
sInT_{upto}(y) &\leftarrow inf(y), out_S(y). \\
sInT_{upto}(y) &\leftarrow succ(z, y), in_S(y), in_T(y), sInT_{upto}(z). \\
sInT_{upto}(y) &\leftarrow succ(z, y), out_S(y), sInT_{upto}(z). \\
sInT &\leftarrow sup(y), sInT_{upto}(y). \\
spoil &\leftarrow sInT.
\end{aligned}$$

We also spoil the answer set if T is not a preferred extension, determined by an external atom with the semantic function $f_{\&argSemExt}$ such that $f_{\&argSemExt}(\mathbf{A}, pref, arg, att, in_T, unused, spoil) = 1$ iff $\mathbf{F}spoil \in \mathbf{A}$ or the extension of predicate in_T is a preferred set extension of the argumentation framework specified by the extension of predicates arg and att . This check is performed using an ASP encoding for preferred extensions from [Egly et al., 2010].

$$\begin{aligned}
tIsNotPref &\leftarrow \&argSemExt[pref, arg, att, in_T, unused, spoil](). \\
spoil &\leftarrow tIsNotPref.
\end{aligned}$$

Note that the parameters $pref$ and $unused$ support more general functionalities of $f_{\&argSemExt}$ which are not relevant for this benchmark. We create a unique answer set whenever $spoil$ is true and require that only spoiled answer sets are returned.

$$\begin{aligned}
in_T(x) &\leftarrow spoil, arg(x). & out_T(x) &\leftarrow spoil, arg(x). \\
sInT &\leftarrow spoil. & tIsNotPref &\leftarrow spoil. & & \leftarrow not\ spoil.
\end{aligned}$$

Given an instance encoded as above, an answer set to the above program exists iff there exists an ideal set extension of the given argumentation framework.

A.2 Conformant Planning

The *Conformant Planning* benchmark results in Section 6.1 were obtained using the following encoding.

Input instances of this benchmark are defined over a set R of robots, sets X and Y of valid x and y coordinates of the environment, and a maximum plan length l ; an instance contains for each robot $r \in R$ the initial position (x, y) as facts $robo_X(r, x, 0)$ and $robo_Y(r, y, 0)$. The encoding consists of the following rules where, unless stated otherwise, $0 \leq t < l$, $r \in R$, $x \in X$, $y \in Y$.

For each robot we generate four possible moves in the environment.

$$\begin{aligned}
move(r, x, y + 1, t) \vee \overline{move}(r, x, y + 1, t) &\leftarrow robo_X(r, x, t), robo_Y(r, y, t). & (y + 1 \in Y) \\
move(r, x, y - 1, t) \vee \overline{move}(r, x, y - 1, t) &\leftarrow robo_X(r, x, t), robo_Y(r, y, t). & (y - 1 \in Y) \\
move(r, x + 1, y, t) \vee \overline{move}(r, x + 1, y, t) &\leftarrow robo_X(r, x, t), robo_Y(r, y, t). & (x + 1 \in X) \\
move(r, x - 1, y, t) \vee \overline{move}(r, x - 1, y, t) &\leftarrow robo_X(r, x, t), robo_Y(r, y, t). & (x - 1 \in X)
\end{aligned}$$

We disallow moving to multiple locations and standing still (the latter is not strictly necessary but we obtained

experimental results that way).

$$\begin{aligned}
&\leftarrow \text{move}(r, x_1, y_1, t), \text{move}(r, x_1, y_2, t). && (x_1, x_2 \in X, x_1 < x_2, y_1, y_2 \in Y) \\
&\leftarrow \text{move}(r, x, y_1, t), \text{move}(r, x, y_2, t). && (y_1, y_2 \in Y, y_1 < y_2) \\
\text{move}_{\exists}(r, t) &\leftarrow \text{move}(r, x, y, t). \\
&\leftarrow \text{not move}_{\exists}(r, t).
\end{aligned}$$

The effect of moving is a deterministic change of location.

$$\text{robo}_X(r, x, t + 1) \leftarrow \text{move}(r, x, y, t). \quad \text{robo}_Y(r, y, t + 1) \leftarrow \text{move}(r, x, y, t).$$

For saturation we guess the position of the object.

$$\text{obj}_X(x) \vee \overline{\text{obj}}_X(x) \leftarrow . \quad \text{obj}_Y(y) \vee \overline{\text{obj}}_Y(y) \leftarrow .$$

We spoil the answer set if the object is at multiple locations.

$$\begin{aligned}
\text{spoil} &\leftarrow \text{obj}_X(x_1), \text{obj}_X(x_2). && (x_1, x_2 \in X, x_1 < x_2) \\
\text{spoil} &\leftarrow \text{obj}_Y(y_1), \text{obj}_Y(y_2). && (y_1, y_2 \in Y, y_1 < y_2)
\end{aligned}$$

We spoil the answer set if the object is at no location.

$$\begin{aligned}
\text{objectHasNoXUpTo}(1) &\leftarrow \overline{\text{obj}}_X(1). \\
\text{objectHasNoXUpTo}(x) &\leftarrow \text{objectHasNoXUpTo}(x - 1), \overline{\text{obj}}_X(x). && (x - 1 \in X) \\
\text{spoil} &\leftarrow \text{objectHasNoXUpTo}(x_{\max}). && (x_{\max} = \max(X)) \\
\text{objectHasNoYUpTo}(1) &\leftarrow \overline{\text{obj}}_Y(1). \\
\text{objectHasNoYUpTo}(y) &\leftarrow \text{objectHasNoYUpTo}(y - 1), \overline{\text{obj}}_Y(y). && (y - 1 \in Y) \\
\text{spoil} &\leftarrow \text{objectHasNoYUpTo}(y_{\max}). && (y_{\max} = \max(Y))
\end{aligned}$$

We spoil the answer set if the object is sensed, which is determined by an external atom with the semantic function $f_{\&sense}$ such that $f_{\&sense}(\mathbf{A}, \text{robo}_X, \text{robo}_Y, \text{obj}_X, \text{obj}_Y, \text{range}, \text{spoil}) = 1$ iff $\mathbf{T}\text{spoil} \in \mathbf{A}$ or the predicates $\text{robo}_X, \text{robo}_Y, \text{obj}_X, \text{obj}_Y$ represent in \mathbf{A} a state where the robot has a distance less than range to the object, i.e., the robot can detect the object.

$$\text{spoil} \leftarrow \&sense[\text{robo}_X, \text{robo}_Y, \text{obj}_X, \text{obj}_Y, \text{range}, \text{spoil}]().$$

We create a unique answer set whenever spoil is true and require that only spoiled answer sets are returned.

$$\begin{aligned}
\text{obj}_X(x) &\leftarrow \text{spoil}. && \overline{\text{obj}}_X(x) \leftarrow \text{spoil}. \\
\text{obj}_Y(y) &\leftarrow \text{spoil}. && \overline{\text{obj}}_Y(y) \leftarrow \text{spoil}. \\
\text{objectHasNoXUpTo}(x) &\leftarrow \text{spoil}. && \text{objectHasNoYUpTo}(y) \leftarrow \text{spoil}. \\
&\leftarrow \text{not spoil}.
\end{aligned}$$

Given an instance encoded as above, an answer set to the above program exists iff there exists a sequence of movements that ensures to detect the object no matter where it is located. Furthermore the movements required to detect the object, i.e., the conformant plan, is encoded in the answer set in the extension of the move predicate.

B Proofs

Proof of Theorem 3.3. The proof is as in [Faber et al., 2011] with external atoms in place of aggregates. \square

Proof of Proposition 3.7. We proceed by contraposition and show that if $I_\Gamma(U, \Gamma_\Pi^{\mathbf{A}}, \Pi, \mathbf{A})$ is not a solution to $\Gamma_\Pi^{\mathbf{A}}$, then U cannot be an unfounded set such that $\mathbf{A}^{\mathbf{T}} \cap U \neq \emptyset$.

First observe that the nogoods in $H_{r,\mathbf{A}}$ demand $\mathbf{T}h_r$ to be true for a rule $r \in \Pi$ in a solution to $\Gamma_\Pi^{\mathbf{A}}$ if and only if some head atom $h \in H(r)$ of this rule is in U . As the truth values of h_r and all $h \in H(r)$ are defined in $I_\Gamma(U, \Gamma_\Pi^{\mathbf{A}}, \Pi, \mathbf{A})$ exactly to this criterion, no nogood from $H_{r,\mathbf{A}}$ can be involved in a contradiction. Furthermore, the nogood $\{\mathbf{F}a \mid \mathbf{T}a \in \mathbf{A}\} \in N_{\Gamma_\Pi^{\mathbf{A}}}$ is violated by $I_\Gamma(U, \Gamma_\Pi^{\mathbf{A}}, \Pi, \mathbf{A})$ only if $\mathbf{A}^{\mathbf{T}} \cap U = \emptyset$; hence, if $I_\Gamma(U, \Gamma_\Pi^{\mathbf{A}}, \Pi, \mathbf{A})$ is not a solution to $\Gamma_\Pi^{\mathbf{A}}$ and $\mathbf{A}^{\mathbf{T}} \cap U \neq \emptyset$, as $O_{\Gamma_\Pi^{\mathbf{A}}}$ is conservative, then for some rule $r \in \Pi$ the nogood in $C_{r,\mathbf{A}}$ must be violated. That is, we know the following: $\mathbf{T}h_r \in I_\Gamma(U, \Gamma_\Pi^{\mathbf{A}}, \Pi, \mathbf{A})$ (and therefore $H(r) \cap U \neq \emptyset$), $\mathbf{F}a \in I_\Gamma(U, \Gamma_\Pi^{\mathbf{A}}, \Pi, \mathbf{A})$ for all $a \in B_o^+(r)$, $\mathbf{t}a \in I_\Gamma(U, \Gamma_\Pi^{\mathbf{A}}, \Pi, \mathbf{A})$ for all $a \in B_e(\hat{r})$, and $\mathbf{T}h \in I_\Gamma(U, \Gamma_\Pi^{\mathbf{A}}, \Pi, \mathbf{A})$ for all $h \in H(r)$ with $\mathbf{A} \models h$. Moreover, we have $C_{r,\mathbf{A}} \neq \emptyset$. We now show that this implies that none of the conditions (i)–(iii) of Definition 3.1 holds for r wrt. U and \mathbf{A} , which means that U is not an unfounded set.

Condition (i) does not hold for r because $\mathbf{A} \models B(r)$ (otherwise $C_{r,\mathbf{A}} = \emptyset$).

Condition (ii) does not hold for r . Suppose to the contrary that it holds. Then there must be some $b \in B(r)$ s.t. $\mathbf{A} \dot{\cup} \neg.U \not\models b$. Because $C_{r,\mathbf{A}} \neq \emptyset$, we know that $\mathbf{A} \models b$. We make a case distinction on the type of b :

- If b is a positive default literal from $A(\Pi)$, then $\mathbf{F}b \in I_\Gamma(U, \Gamma_\Pi^{\mathbf{A}}, \Pi, \mathbf{A})$ and therefore $b \notin U$. Consequently $\mathbf{A} \dot{\cup} \neg.U \models b$. Contradiction.
- If b is a negative default literal over $A(\Pi)$, then $\mathbf{A} \models b$ implies $\mathbf{A} \dot{\cup} \neg.U \models b$. Contradiction.
- If b is a positive or default-negated replacement atom, then $\mathbf{t}b \in I_\Gamma(U, \Gamma_\Pi^{\mathbf{A}}, \Pi, \mathbf{A})$. But this implies, by definition of $I_\Gamma(U, \Gamma_\Pi^{\mathbf{A}}, \Pi, \mathbf{A})$, that $\mathbf{A} \dot{\cup} \neg.U \models b$. Contradiction.

Condition (iii) does not hold for r because $\mathbf{T}h \in I_\Gamma(U, \Gamma_\Pi^{\mathbf{A}}, \Pi, \mathbf{A})$ and thus, by definition of $I_\Gamma(U, \Gamma_\Pi^{\mathbf{A}}, \Pi, \mathbf{A})$, $h \in U$ for all $h \in H(r)$ with $\mathbf{A} \models h$. Thus $\mathbf{A} \not\models a$ for all $a \in H(r) \setminus U$. \square

Proof of Theorem 3.9. Suppose U is not an unfounded set. Then there is an $r \in \Pi$ s.t. $H(r) \cap U \neq \emptyset$ and none of the conditions (i)–(iii) in Definition 3.1 is satisfied. We show now that S cannot be a solution to $\Gamma_\Pi^{\mathbf{A}}$ that satisfies conditions (a) and (b), which proves the result.

Because condition (i) does not hold, there is a nogood of form

$$\{\{\mathbf{T}h_r\} \cup \{\mathbf{F}a \mid a \in B_o^+(r), \mathbf{A} \models a\} \cup \{\mathbf{t}a \mid a \in B_e(\hat{r})\} \cup \{\mathbf{T}h \mid h \in H(r), \mathbf{A} \models h\}\}$$

in $\Gamma_\Pi^{\mathbf{A}}$.

We now show that S contains all signed literals of this nogood, i.e., the nogood is violated by the assignment S .

Since $H(r) \cap U \neq \emptyset$, $\mathbf{T}h_r \in S$ (otherwise a nogood in $H_r^{\mathbf{A}}$ is violated). As U is not an unfounded set, Condition (ii) in Definition 3.1 does not hold. Consider any $a \in B_o^+(r)$ s.t. $\mathbf{A} \models a$. Then $a \notin U$, otherwise $A \dot{\cup} \neg.U \not\models a$ and we have a contradiction with the assumption that Condition (ii) is unsatisfied. But then $\mathbf{F}a \in S$ (because S is complete and would imply $a \in U$ otherwise).

Now consider any $\&g[\mathbf{p}](\mathbf{c}) \in B_e(r)$. Then $\mathbf{A} \dot{\cup} \neg.U \models \&g[\mathbf{p}](\mathbf{c})$ (as (ii) is violated). If $\mathbf{A} \not\models \&g[\mathbf{p}](\mathbf{c})$, then Condition (i) would be satisfied, hence $\mathbf{A} \models \&g[\mathbf{p}](\mathbf{c})$. But then $\mathbf{T}e_{\&g[\mathbf{p}]}(\mathbf{c}) \in S$, otherwise $\mathbf{A} \dot{\cup} \neg.U \not\models \&g[\mathbf{p}](\mathbf{c})$ by precondition (b) of this proposition. Next consider all *not* $\&g[\mathbf{p}](\mathbf{c}) \in B_e(r)$. Then

$\mathbf{A} \dot{\cup} \neg.U \not\models \&g[\mathbf{p}](\mathbf{c})$ (as (ii) is violated). If $\mathbf{A} \models \&g[\mathbf{p}](\mathbf{c})$, then Condition (i) would be satisfied, hence $\mathbf{A} \not\models \&g[\mathbf{p}](\mathbf{c})$. But then $\mathbf{F}e_{\&g[\mathbf{p}]}(\mathbf{c}) \in S$, otherwise $\mathbf{A} \dot{\cup} \neg.U \models \&g[\mathbf{p}](\mathbf{c})$ by precondition (a) of this proposition. Therefore, we have $\mathbf{t}a \in S$ for all $a \in B_e(\hat{r})$.

Finally, because Condition (iii) in Definition 3.1 does not hold, $h \in U$ and therefore also $\mathbf{T}h \in S$ for all $h \in H(r)$ with $\mathbf{A} \models a$.

This concludes the proof that S cannot be a solution to $\Gamma_{\Pi}^{\mathbf{A}}$ satisfying (a) and (b), if U is not an unfounded set. \square

Proof of Proposition 3.13. We proceed by contraposition and show that if $I_{\Omega}(U, \Omega_{\Pi}, \Pi, \mathbf{A})$ is not a solution to Ω_{Π} with assumptions $\mathcal{A}_{\mathbf{A}}$, then U cannot be an unfounded set such that $\mathbf{A}^{\mathbf{T}} \cap U \neq \emptyset$.

First observe that the nogoods in H_r demand $\mathbf{T}h_r$ to be true for a rule $r \in \Pi$ if and only if some head atom $h \in H(r)$ of this rule is in U . Moreover, the nogoods in D_a for each $a \in A(\Pi)$ force $a_{\mathbf{A} \dot{\cup} \neg.U}$ to true if and only if $\mathbf{T}a \in \mathbf{A}$ and $a \notin U$, which is equivalent to $\mathbf{T}a \in \mathbf{A} \dot{\cup} \neg.U$; $a_{\mathbf{A} \wedge U}$ to true if and only if $\mathbf{T}a \in \mathbf{A}$ and $a \in U$; and $a_{\overline{\mathbf{A}} \vee U}$ to true if and only if $\mathbf{F}a \in \mathbf{A}$ or $a \in U$. As the truth values of h_r for each $r \in \Pi$, and $a_{\mathbf{A} \dot{\cup} \neg.U}$ and $a_{\mathbf{A} \wedge U}$ and $a_{\overline{\mathbf{A}} \vee U}$ for each $a \in A(\Pi)$ in $I_{\Omega}(U, \Omega_{\Pi}, \Pi, \mathbf{A})$ are defined exactly to these conditions, a contradiction must involve C_r for some $r \in \Pi$. Furthermore, the nogood $\{\mathbf{F}a \mid a \in A(\Pi)\}$ is violated by $I_{\Omega}(U, \Omega_{\Pi}, \Pi, \mathbf{A})$ only if $A(\Pi) \cap U = \emptyset$, and thus $\mathbf{A}^{\mathbf{T}} \cap U \neq \emptyset$; hence, if $I_{\Omega}(U, \Omega_{\Pi}, \Pi, \mathbf{A})$ is not a solution to Ω_{Π} such that $\mathbf{A}^{\mathbf{T}} \cap U \neq \emptyset$, since $O_{\Omega, \Pi}$ is conservative, for some rule $r \in \Pi$ the nogood in C_r must be violated. That is, we know the following:

- (I) $\mathbf{T}h_r \in I_{\Omega}(U, \Omega_{\Pi}, \Pi, \mathbf{A})$ (and therefore $H(r) \cap U \neq \emptyset$),
- (II) $\mathbf{T}a_{\mathbf{A}} \in I_{\Omega}(U, \Omega_{\Pi}, \Pi, \mathbf{A})$ for all $a \in B^+(\hat{r})$ and $\mathbf{F}a_{\mathbf{A}} \in I_{\Omega}(U, \Omega_{\Pi}, \Pi, \mathbf{A})$ for all $a \in B^-(\hat{r})$,
- (III) $\mathbf{F}a_{\mathbf{A} \wedge U} \in I_{\Omega}(U, \Omega_{\Pi}, \Pi, \mathbf{A})$ for all $a \in B_o^+(r)$, $\mathbf{t}a \in I_{\Omega}(U, \Omega_{\Pi}, \Pi, \mathbf{A})$ for all $a \in B_e(\hat{r})$, and
- (IV) $\mathbf{T}h_{\overline{\mathbf{A}} \vee U} \in I_{\Omega}(U, \Omega_{\Pi}, \Pi, \mathbf{A})$ for all $h \in H(r)$.

We now show that this implies that none of the conditions of Definition 3.1 holds for r wrt. U and \mathbf{A} , which means that U is not an unfounded set (h_r is true in $I_{\Omega}(U, \Omega_{\Pi}, \Pi, \mathbf{A})$, which implies $H(r) \cap U \neq \emptyset$).

Condition (i) does not hold for r because of (II), which by our assumptions $\mathcal{A}_{\mathbf{A}}$ implies $\mathbf{A} \models B(r)$. Condition (ii) does not hold for r . Suppose to the contrary that it holds. Then there must be some $b \in B(r)$ s.t. $\mathbf{A} \dot{\cup} \neg.U \not\models b$. Since condition (i) is already known to be violated, we can assume that $\mathbf{A} \models b$. We make a case distinction on the type of b :

- If b is a positive default literal from $A(\Pi)$, then $b \in U$ (otherwise $\mathbf{A} \dot{\cup} \neg.U \models b$). But then we have by definition of $I_{\Omega}(U, \Omega_{\Pi}, \Pi, \mathbf{A})$ that $\mathbf{T}b_{\mathbf{A} \wedge U} \in I_{\Omega}(U, \Omega_{\Pi}, \Pi, \mathbf{A})$, which contradicts (III).
- If b is a negative default literal over $A(\Pi)$, then $\mathbf{A} \models b$ implies $\mathbf{A} \dot{\cup} \neg.U \models b$. Contradiction.
- If b is a positive or default-negated replacement atom of an external, then $\mathbf{t}b \in I_{\Omega}(U, \Omega_{\Pi}, \Pi, \mathbf{A})$. But this implies, by definition of $I_{\Omega}(U, \Omega_{\Pi}, \Pi, \mathbf{A})$, that $\mathbf{A} \dot{\cup} \neg.U \models b$. Contradiction.

Condition (iii) does not hold for r because $\mathbf{T}h_{\overline{\mathbf{A}} \vee U} \in I_{\Omega}(U, \Omega_{\Pi}, \Pi, \mathbf{A})$ and thus, by definition of $I_{\Omega}(U, \Omega_{\Pi}, \Pi, \mathbf{A})$, $h \in U$ for all $h \in H(r)$ with $\mathbf{A} \models h$. Thus $\mathbf{A} \not\models a$ for all $a \in H(r) \setminus U$. \square

Proof of Theorem 3.15. Suppose U is not an unfounded set. Then some $r \in \Pi$ exists such that $H(r) \cap U \neq \emptyset$ and none of the conditions (i)–(iii) in Definition 3.1 is satisfied. We show that then S , assuming that $\mathcal{A}_{\mathbf{A}}$ is satisfied and (a) and (b) hold, cannot be a solution to Ω_{Π} .

Due to rule r , Ω_{Π} (more specifically, $N_{\Omega, \Pi}$) contains a nogood N of form

$$\begin{aligned}
N = & \{ \{ \mathbf{T}h_r \} \cup \\
& \{ \mathbf{T}a_{\mathbf{A}} \mid a \in B^+(\hat{r}) \} \cup \{ \mathbf{F}a_{\mathbf{A}} \mid a \in B^-(\hat{r}) \} \cup \\
& \{ \mathbf{F}a_{\mathbf{A} \wedge U} \mid a \in B_o^+(r) \} \cup \{ \mathbf{t}a \mid a \in B_e(\hat{r}) \} \cup \\
& \{ \mathbf{T}h_{\overline{\mathbf{A}} \vee U} \mid h \in H(r) \} \}.
\end{aligned}$$

We show that S contains all signed literals of N , i.e., N is violated by S . As $H(r) \cap U \neq \emptyset$, $\mathbf{T}h_r \in S$ (otherwise a nogood in H_r is violated). Furthermore, as U is not an unfounded set, condition (i) in Definition 3.1 does not hold for U wrt. \mathbf{A} , and hence $\mathbf{A} \models B(r)$. The assumptions $\mathcal{A}_{\mathbf{A}}$ thus enforce that $\mathbf{T}a_{\mathbf{A}} \in S$ for all $a \in B^+(\hat{r})$ and $\mathbf{F}a_{\mathbf{A}} \in S$ for all $a \in B^-(\hat{r})$.

Consider next an arbitrary $a \in B_o^+(r)$. As condition (ii) in Definition 3.1 does not hold for U wrt. \mathbf{A} , we have $\mathbf{A} \models a$ and $a \notin U$; the latter implies by definition of U that $\mathbf{F}a \in S$. From the nogood $\{ \mathbf{T}a_{\mathbf{A} \wedge U}, \mathbf{F}a \}$ we conclude $\mathbf{F}a_{\mathbf{A} \wedge U} \in S$.

We next show that for every $a \in B_e(\hat{r})$, it holds that $\mathbf{t}a \in S$. Indeed, let $\&g[\mathbf{p}](\mathbf{c}) \in B_e(r)$. We have $\mathbf{A} \cup \neg.U \models \&g[\mathbf{p}](\mathbf{c})$ (as condition (ii) is violated). Furthermore, $\mathbf{A} \models \&g[\mathbf{p}](\mathbf{c})$, as $\mathbf{A} \not\models \&g[\mathbf{p}](\mathbf{c})$ would imply that condition (i) is satisfied. Thus from condition (b) of the hypothesis, it follows that $\mathbf{T}e_{\&g[\mathbf{p}](\mathbf{c})} \in S$. Similarly, let $\text{not } \&g[\mathbf{p}](\mathbf{c}) \in B_e(r)$. Then $\mathbf{A} \cup \neg.U \not\models \&g[\mathbf{p}](\mathbf{c})$ (as condition (ii) is violated) and $\mathbf{A} \not\models \&g[\mathbf{p}](\mathbf{c})$ as $\mathbf{A} \models \&g[\mathbf{p}](\mathbf{c})$ would satisfy condition (i). Thus from condition (a) of the hypothesis, it follows that $\mathbf{F}e_{\&g[\mathbf{p}](\mathbf{c})} \in S$. Thus we have $\mathbf{t}a \in S$ for all $a \in B_e(\hat{r})$.

Finally, because condition (iii) in Definition 3.1 does not hold for U wrt. \mathbf{A} , we have $h \in U$ and therefore also $\mathbf{T}h \in S$ for all $h \in H(r)$ with $\mathbf{A} \models a$. That is, for each $h \in H(r)$, either $\mathbf{F}h_{\mathbf{A}} \in S$ or $\mathbf{T}h \in S$. But by the nogoods $\{ \mathbf{F}h_{\overline{\mathbf{A}} \vee U}, \mathbf{F}h_{\mathbf{A}} \}, \{ \mathbf{F}h_{\overline{\mathbf{A}} \vee U}, \mathbf{T}h \} \in D_h$, in both cases we have $\mathbf{T}h_{\overline{\mathbf{A}} \vee U} \in S$.

Hence, $N \subseteq S$ holds, i.e., N is violated by S , and thus S is not a solution of Ω_{Π} . This completes the proof of the result. \square

Proof of Proposition 4.1. To show that $U \setminus \{a\}$ is an unfounded set wrt. \mathbf{A} , consider $r \in \Pi$ such that $H(r) \cap (U \setminus \{a\}) \neq \emptyset$. We have to show that one of the conditions (i)–(iii) of Definition 3.1 holds for r . By hypothesis, U is an unfounded set of Π wrt. \mathbf{A} , and $H(r) \cap (U \setminus \{a\}) \neq \emptyset$ implies $H(r) \cap U \neq \emptyset$. If condition (i) holds for U , it also holds wrt. $U \setminus \{a\}$ because this condition depends only on r and \mathbf{A} . Also if condition (ii) holds for U , it also holds wrt. $U \setminus \{a\}$ because $\mathbf{A} \cup \neg.U$ is equivalent to $\mathbf{A} \cup \neg.(U \setminus \{a\})$ since $\mathbf{A} \not\models a$. Finally, if condition (iii) holds for U , then some atom $b \in H(r) \setminus U$ exists such that $\mathbf{A} \models b$. As $\mathbf{A} \not\models a$, it follows $a \neq b$ and hence condition (iii) holds for $U \setminus \{a\}$. This proves the result. \square

Proof of Proposition 4.2. Suppose that changing the truth value of b in S turns the solution to a counterexample S_b of $\Gamma_{\Pi}^{\mathbf{A}}$ (resp. Ω_{Π}). That is, S_b must violate some nogood $N \in \Gamma_{\Pi}^{\mathbf{A}}$ resp. ($N \in \Omega_{\Pi}$) containing b , i.e., either $\mathbf{T}b \in N$ or $\mathbf{F}b \in N$.

For the encoding $\Gamma_{\Pi}^{\mathbf{A}}$, the nogood N corresponds to a rule with $b \in B^+(\hat{r})$ or $b \in B^-(\hat{r})$ and $\mathbf{A} \models B(r)$, and N contains also the signed literals (1) $\mathbf{F}a$ for all $a \in B_o^+$ with $\mathbf{A} \models a$ and (2) $\mathbf{T}a$ for all $a \in H(r)$ with $\mathbf{A} \models a$. By hypothesis, we have either (a) $\mathbf{T}a \in S$ for some $a \in B_o^+(r)$ with $\mathbf{A} \models a$, or (b) $\mathbf{F}a$ for some $a \in H(r)$ with $\mathbf{A} \models a$. But then the nogood cannot be violated, because (a) contradicts one of (1) and (b) contradicts one of (2).

For the encoding Ω_{Π} , the nogood N corresponds to a rule r with $b \in B^+(\hat{r})$ or $b \in B^-(\hat{r})$. The nogood contains also the signed literals (1) $\mathbf{T}a_{\mathbf{A}}$ for all $a \in B^+(\hat{r})$ and $\mathbf{F}a_{\mathbf{A}}$ for all $a \in B^-(\hat{r})$, (2) $\mathbf{F}a_{\mathbf{A} \wedge U}$ for all $a \in B_o^+$, and (3) $\mathbf{T}h_{\overline{\mathbf{A}} \vee U}$ for all $h \in H(r)$. Because of (1) and since \mathbf{A} is a solution to $\mathcal{A}_{\mathbf{A}}$, we have $\mathbf{A} \models B(r)$. By hypothesis, we have either (a) $\mathbf{T}a \in S$ for some $a \in B_o^+(r)$ with $\mathbf{A} \models a$, or (b) $\mathbf{F}a$ for some $a \in H(r)$ with $\mathbf{A} \models a$. But then the nogood N cannot be violated, because (a) contradicts part (2) by definition of $\mathcal{A}_{\mathbf{A}}$ and $a_{\mathbf{A} \wedge U}$, and (b) contradicts part (3) by definition of $\mathcal{A}_{\mathbf{A}}$ and $h_{\overline{\mathbf{A}} \vee U}$. \square

Proof of Proposition 4.6. If $\mathcal{T}_\Gamma(N, \mathbf{A}) = \emptyset$ then the proposition trivially holds. Otherwise $\mathcal{T}_\Gamma(N, \mathbf{A}) = \{C\}$ and we know that $\mathbf{T}t_i \in \mathbf{A}$ for all $1 \leq i \leq n$. Suppose C is violated by $I_\Gamma(U, \Gamma_\Pi^{\mathbf{A}}, \Pi, \mathbf{A})$. Then $\mathbf{F}t_i \in I_\Gamma(U, \Gamma_\Pi^{\mathbf{A}}, \Pi, \mathbf{A})$ and therefore $t_i \notin U$ for all $1 \leq i \leq n$, and $\mathbf{T}f_i \in I_\Gamma(U, \Gamma_\Pi^{\mathbf{A}}, \Pi, \mathbf{A})$ for all $1 \leq i \leq m$ with $\mathbf{A} \models f_i$, and $\sigma e_{\&g[\mathbf{p}]}(\mathbf{c}) \in I_\Gamma(U, \Gamma_\Pi^{\mathbf{A}}, \Pi, \mathbf{A})$.

But then $\mathbf{A} \dot{\cup} \neg.U \models t_i$ for all $1 \leq i \leq n$ and $\mathbf{A} \dot{\cup} \neg.U \not\models f_i$ for all $1 \leq i \leq m$. Because the nogood N is a valid input-output-relationship, this implies $\bar{\sigma}e_{\&g[\mathbf{p}]}(\mathbf{c}) \in \mathbf{A} \dot{\cup} \neg.U$. By definition of $I_\Gamma(U, \Gamma_\Pi^{\mathbf{A}}, \Pi, \mathbf{A})$, we then have $\bar{\sigma}e_{\&g[\mathbf{p}]}(\mathbf{c}) \in I_\Gamma(U, \Gamma_\Pi^{\mathbf{A}}, \Pi, \mathbf{A})$, which contradicts the assumption that $\mathcal{T}_\Gamma(N, \mathbf{A})$ is violated. \square

Proof of Proposition 4.9. We know $\mathcal{T}_\Omega(N) = \{C\}$. Suppose $I_\Omega(U, \Omega_\Pi, \Pi, \mathbf{A})$ violates C . Then $\mathbf{T}t_{i\mathbf{A}} \in I_\Omega(U, \Omega_\Pi, \Pi, \mathbf{A})$ and therefore $\mathbf{T}t_i \in \mathbf{A}$ and $\mathbf{F}t_i \in I_\Omega(U, \Omega_\Pi, \Pi, \mathbf{A})$ for all $1 \leq i \leq n$; $\mathbf{F}f_{i\mathbf{A} \dot{\cup} \neg.U} \in I_\Omega(U, \Omega_\Pi, \Pi, \mathbf{A})$ and therefore either $\mathbf{F}f_i \in \mathbf{A}$ or $f_i \in U$, for all $1 \leq i \leq m$; and $\sigma e_{\&g[\mathbf{p}]}(\mathbf{c}) \in I(U, \Omega_\Pi, \mathbf{A})$.

But then, by definition of $I_\Omega(U, \Omega_\Pi, \Pi, \mathbf{A})$, $\mathbf{T}t_i \in \mathbf{A}$ and $t_i \notin U$ for all $1 \leq i \leq n$, hence $\mathbf{A} \dot{\cup} \neg.U \models t_i$ for all $1 \leq i \leq n$. Moreover, $\mathbf{A} \dot{\cup} \neg.U \not\models f_i$ for all $1 \leq i \leq m$. Because nogood N is a valid input-output-relationship, this implies $\bar{\sigma}e_{\&g[\mathbf{p}]}(\mathbf{c}) \in \mathbf{A} \dot{\cup} \neg.U$. On the other hand, by definition of $I_\Omega(U, \Omega_\Pi, \Pi, \mathbf{A})$, we have $\bar{\sigma}e_{\&g[\mathbf{p}]}(\mathbf{c}) \in I_\Omega(U, \Omega_\Pi, \Pi, \mathbf{A})$; this contradicts the assumption that the nogood C is violated. \square

Proof of Proposition 4.12. Suppose some answer set \mathbf{A}' of Π exists which is not a solution to $L_1(U, \Pi, \mathbf{A})$, i.e., it violates some nogood $N = \{\sigma_0, \sigma_1, \dots, \sigma_n\} \in L_1(U, \Pi, \mathbf{A})$. We show that in this case U is an unfounded set of Π wrt. \mathbf{A}' such that $U \cap \mathbf{A}' \neq \emptyset$; this means that \mathbf{A}' is not-unfounded-free, which contradicts that \mathbf{A}' is an answer set of Π .

Let $r \in \Pi$ be a rule such that $H(r) \cap U \neq \emptyset$. We have to show that one of the conditions (i)–(iii) of Definition 3.1 holds.

If $B_o^+(r) \cap U \neq \emptyset$, then condition (ii) holds. Hence we assume in the following that $B_o^+(r) \cap U = \emptyset$, which means that r is an external rule of Π wrt. U . But then for some $\sigma_i \in N$ with $1 \leq i \leq n$ we have either (1) $\sigma_i = \mathbf{T}h$ for some $h \in H(r)$ with $h \notin U$ and $\mathbf{A} \models h$, or (2) $\sigma_i = \mathbf{F}b$ for some $b \in B_o^+(r)$ with $\mathbf{A} \not\models b$. Since \mathbf{A}' violates N by assumption, we have $\sigma_i \in \mathbf{A}'$. In case (1) the condition (iii) is satisfied, while in case (2) then condition (i) is satisfied.

Moreover, by definition of L_1 some $a \in U$ exists such that $\mathbf{T}a \in \mathbf{A}'$, i.e., \mathbf{A}' intersects with U . This proves the result. \square

Proof of Proposition 4.15. Towards a contradiction, suppose some answer set \mathbf{A}' of Π is not a solution to $L_2(U, \Pi, \mathbf{A})$. Let $N = \{\mathbf{T}a \mid a \in \mathbf{A} \dot{\cup} \neg.U\} \cup \{\sigma_0, \sigma_1, \dots, \sigma_n\} \in L_2(U, \Pi, \mathbf{A})$ be a violated nogood. Because $\sigma_i \in \mathbf{A}'$ for all $1 \leq i \leq n$, we know that \mathbf{A}' falsifies (at least) the bodies of rules in Π that are falsified by \mathbf{A} ; consequently, $f\Pi^{\mathbf{A}'} \subseteq f\Pi^{\mathbf{A}}$. From $\mathbf{A} \models \Pi$ and the hypothesis that U is an unfounded set it follows that $\mathbf{A} \dot{\cup} \neg.U \models f\Pi^{\mathbf{A}}$; hence, also $\mathbf{A} \dot{\cup} \neg.U \models f\Pi^{\mathbf{A}'}$. Moreover, $\mathbf{T}a \in \mathbf{A}'$ for all $a \in \mathbf{A} \dot{\cup} \neg.U$, and therefore $\mathbf{A}'^{\mathbf{T}} \supseteq (\mathbf{A} \dot{\cup} \neg.U)^{\mathbf{T}}$. Because $\sigma_0 \in \mathbf{A}'$, we conclude $\mathbf{A}'^{\mathbf{T}} \supsetneq (\mathbf{A} \dot{\cup} \neg.U)^{\mathbf{T}}$, i.e., \mathbf{A}' is not a subset-minimal model of $\Pi^{\mathbf{A}'}$. Consequently, \mathbf{A}' is not an answer set of Π , which is a contradiction. \square

Proof of Lemma 5.3. If $Y = \emptyset$, then the result holds trivially. Otherwise, let $r \in \Pi$ with $H(r) \cap Y \neq \emptyset$. We show that one of the conditions (i)–(iii) in Definition 3.1 holds. Observe that $H(r) \cap U \neq \emptyset$ because $U \supseteq Y$. Since U is an unfounded set of Π wrt. \mathbf{A} , either

- (i) $\mathbf{A} \not\models b$ for some $b \in B(r)$; or
- (ii) $\mathbf{A} \dot{\cup} \neg.U \not\models b$ for some $b \in B(r)$; or
- (iii) $\mathbf{A} \models h$ for some $h \in H(r) \setminus U$.

In case (i), the condition also holds wrt. Y . In case (ii), let $a \in H(r)$ such that $a \in Y$, and $b \in B(r)$ such that $\mathbf{A} \dot{\cup} \neg.U \not\models b$. We make a case distinction: either b is an ordinary literal or an external one.

If b is an ordinary default-negated atom *not* c , then $\mathbf{A} \dot{\cup} \neg.U \not\models b$ implies $\mathbf{T}c \in \mathbf{A}$ and $c \notin U$, and therefore also $\mathbf{A} \dot{\cup} \neg.Y \not\models b$. So assume b is an ordinary atom. If $b \notin U$ then $\mathbf{A} \not\models b$ and case (i) applies, so assume $b \in U$. Because $a \in H(r)$ and $b \in B(r)$, we have $a \rightarrow b$ and therefore either $a, b \in C$ or $a, b \in Y$ (because there are no ordinary edges between C and Y). But by assumption $a \in Y$, and therefore $b \in Y$, hence $\mathbf{A} \dot{\cup} \neg.Y \not\models b$.

If b is an external literal, then there is no $q \in U$ with $a \rightarrow_e q$ and $q \notin Y$. Otherwise, this would imply $q \in C$ and C would have an incoming e-edge, which contradicts the assumption that C is a cut. Hence, for all $q \in U$ with $a \rightarrow_e q$, also $q \in Y$, and therefore the truth value of b under $\mathbf{A} \dot{\cup} \neg.U$ and $\mathbf{A} \dot{\cup} \neg.Y$ is the same. Hence $\mathbf{A} \dot{\cup} \neg.Y \not\models b$.

In case (iii), also $\mathbf{A} \models h$ for some $h \in H(r) \setminus Y$ because $Y \subseteq U$ and therefore $H(r) \setminus Y \supseteq H(r) \setminus U$. \square

Proof of Lemma 5.5. If $U = \emptyset$, then the result holds trivially. Otherwise, suppose $\hat{r} \in \hat{\Pi}$ and $a \in H(\hat{r}) \cap U$. Observe that \hat{r} cannot be an external atom guessing rule because U contains only ordinary atoms. We show that one of the conditions in Definition 3.1 holds for \hat{r} wrt. $\hat{\mathbf{A}}$.

Because \hat{r} is no external atom guessing rule, there is a corresponding rule $r \in \Pi$ containing external atoms in place of replacement atoms. Because U is an unfounded set of Π and $H(r) = H(\hat{r})$, either:

- (i) $\mathbf{A} \not\models b$ for some $b \in B(r)$; or
- (ii) $\mathbf{A} \dot{\cup} \neg.U \not\models b$ for some $b \in B(r)$; or
- (iii) $\mathbf{A} \models h$ for some $h \in H(r) \setminus U$

In case (i), let $b \in B(r)$ such that $\mathbf{A} \not\models b$ and \hat{b} the corresponding literal in $B(\hat{b})$ (which is the same if b is ordinary and the corresponding replacement literal if b is external). Then also $\hat{\mathbf{A}} \not\models \hat{b}$ because $\hat{\mathbf{A}}$ is compatible.

In case (ii), we make a case distinction: either b is ordinary or external.

If b is ordinary, then $b \in B(\hat{r})$ and $\hat{\mathbf{A}} \dot{\cup} \neg.U \not\models b$ holds because \mathbf{A} and $\hat{\mathbf{A}}$ are equivalent for ordinary atoms.

If b is an external atom or default-negated external atom, then no atom $p(\mathbf{c}) \in U$ is input to it, i.e. p is not a predicate input parameter of b ; otherwise we had $a \rightarrow_e p(\mathbf{c})$, contradicting our assumption that U has no internal e-edges. But then $\mathbf{A} \dot{\cup} \neg.U$ implies $\mathbf{A} \not\models b$ because the truth value of b under $\mathbf{A} \dot{\cup} \neg.U$ and \mathbf{A} is the same. Therefore we can apply case (i).

In case (iii), also $\hat{\mathbf{A}} \models h$ for some $h \in H(\hat{r}) \setminus U$ because $H(r) = H(\hat{r})$ contains only ordinary atoms and \mathbf{A} is equivalent to $\hat{\mathbf{A}}$ for ordinary atoms. \square

Proof of Theorem 5.7. We define the *reachable set* $R(a)$ from some atom a as

$$R(a) = \{b \mid (a, b) \in \{\rightarrow \cup \leftarrow\}^*\},$$

i.e. the set of atoms $b \in U$ reachable from a using edges from $\rightarrow \cup \leftarrow$ only but no e-edges.

We first assume that U contains at least one e-edge, i.e. there are $x, y \in U$ such that $x \rightarrow_e y$. Now we show that there is a $u \in U$ with outgoing e-edge (i.e. $u \rightarrow_e v$ for some $v \in U$), but such that $R(u)$ has no incoming e-edges from U (i.e. for all $v \in R(u)$ and $b \in U$, $b \not\rightarrow_e v$ holds). Suppose to the contrary that for all a with outgoing e-edges, the reachable set $R(a)$ has an incoming e-edge from U . We now construct an e-cycle

under \rightarrow^d , which contradicts our assumption. Start with an arbitrary node $c_0 \in U$ that has an outgoing e-edge, let p_0 be the (possibly empty) path (under $\rightarrow \cup \leftarrow$) from c_0 to the node $d_0 \in R(c_0)$ such that d_0 has an incoming e-edge, i.e. there is a c_1 such that $c_1 \rightarrow_e d_0$; note that $c_1 \notin R(c_0)$.³ By assumption, also some node d_1 in $R(c_1)$ has an incoming e-edge (from some node $c_2 \notin R(c_1)$). Let p_1 be the path from c_1 to d_1 , etc. By iteration we can construct the concatenation of the paths $p_0, (d_0, c_1), p_1, (d_1, c_2), p_2, \dots, p_i, (d_i, c_{i+1}), \dots$, where the p_i from c_i to d_i are the paths within reachable sets, and the (d_i, c_{i+1}) are the e-edges between reachable sets. However, as U is finite some nodes on this path must be equal, i.e., a subsequence of the constructed sequence represents an e-cycle (in reverse order).

This proves that u is a node with outgoing e-edge but such that $R(u)$ has no incoming e-edges. We next show that $R(u)$ is a cut. Condition (i) is immediately satisfied by definition of u . Condition (ii) is shown as follows. Let $u' \in R(u)$ and $v' \in U \setminus R(u)$. We have to show that $u' \not\rightarrow v'$ and $v' \not\rightarrow u'$. Suppose, towards a contradiction, that $u' \rightarrow v'$. Because $u' \in R(u)$, there is a path from u to u' under $\rightarrow \cup \leftarrow$. But if $u' \rightarrow v'$, then there would also be a path from u to v' under $\rightarrow \cup \leftarrow$ and v' would be in $R(u)$, a contradiction. Analogously, $v' \rightarrow u'$ would also imply that there is a path from u to v' because there is a path from u to u' , again a contradiction.

Therefore, $R(u)$ is a cut of U , and by Lemma 5.3, it follows that $U \setminus R(u)$ is an unfounded set. Observe that $U \setminus R(u)$ contains one e-edge less than U because u has an outgoing e-edge and that $U \setminus R(u) \neq \emptyset$; indeed, by assumption some $w \in U$ exists such that $u \rightarrow_e w$, and clearly $w \notin R(u)$. By iterating this argument, the number of e-edges in the unfounded set can be reduced to zero in a nonempty core. Eventually Lemma 5.5 applies, proving that the remaining set is an unfounded set of $\hat{\Pi}$. \square

Proof of Corollary 5.8. Towards a contradiction, suppose an unfounded set U of Π wrt. \mathbf{A} exists. Then U contains no e-cycle because there is no e-cycle under \rightarrow^d . By Theorem 5.7 there is an unfounded set of $\hat{\Pi}$ wrt. $\hat{\mathbf{A}}$, which contradicts our assumption that $\hat{\Pi}$ has no unfounded set wrt. $\hat{\mathbf{A}}$. \square

Proof of Theorem 5.11. If U contains no cyclic input atoms, then all cycles under \rightarrow^d containing e-edges in the atom dependency graph of Π are broken, i.e. U does not contain an e-cycle under \rightarrow^d . Then by Theorem 5.7 there exists a nonempty unfounded set of $\hat{\Pi}$ wrt. $\hat{\mathbf{A}}$. \square

Proof of Theorem 5.13. Let U be a nonempty unfounded set of Π wrt. \mathbf{A} . Because \mathcal{C} is a decomposition of $A(\Pi)$ into strongly connected components, the *component dependency graph*

$$\langle \mathcal{C}, \{(C_1, C_2) \mid C_1, C_2 \in \mathcal{C}, \exists a_1 \in C_1, a_2 \in C_2 : (a_1, a_2) \in \rightarrow \cup \rightarrow_e\} \rangle$$

is acyclic. Following the hierarchical component dependency graph from the nodes without predecessor components downwards, we can find a “first” component which has a nonempty intersection with U , i.e., there exists a component $C \in \mathcal{C}$ such that $C \cap U \neq \emptyset$ but $C' \cap U = \emptyset$ for all transitive predecessor components C' of C .

We show that $U \cap C$ is an unfounded set of Π_C wrt. \mathbf{A} . Let $r \in \Pi_C$ be a rule such that $H(r) \cap (U \cap C) \neq \emptyset$. We have to show that one of the conditions (i)–(iii) of Definition 3.1 holds for r wrt. \mathbf{A} and $U \cap C$.

Because U is an unfounded set of Π wrt. \mathbf{A} we know that one of the conditions (i)–(iii) holds for r wrt. \mathbf{A} and U . In case of condition (i), then it trivially holds also wrt. \mathbf{A} and $U \cap C$ because these conditions depend only on the assignment \mathbf{A} , but not on the unfounded set U ; in case of condition (iii), it clearly holds because $H(r) \setminus U \neq \emptyset$ is included in $H(r) \setminus (U \cap C)$.

³ Whenever $x \rightarrow_e y$ for $x, y \in U$, then there is no path from x to y under $\rightarrow \cup \leftarrow$, because otherwise we would have an e-cycle under \rightarrow^d .

In case of condition (ii), we have that $\mathbf{A} \dot{\cup} \neg.U \not\models b$ for some (ordinary or external) body literal $b \in B(r)$. We show next that the truth value of all literals in $B(r)$ is the same under $\mathbf{A} \dot{\cup} \neg.U$ and $\mathbf{A} \dot{\cup} \neg.(U \cap C)$, which proves that condition (ii) holds also wrt. \mathbf{A} and $U \cap C$.

If $b = \text{not } a$ for some ordinary atom a , then $\mathbf{T}a \in \mathbf{A}$ and $a \notin U$ and consequently $a \notin U \cap C$, hence $\mathbf{A} \dot{\cup} \neg.(U \cap C) \not\models b$. If b is an ordinary atom, then either $\mathbf{F}b \in \mathbf{A}$, which implies immediatly that $\mathbf{A} \dot{\cup} \neg.(U \cap C) \not\models b$, or $b \in U$. But in the latter case b is either in a predecessor component C' of C or in C itself (since $h \rightarrow b$ for all $h \in H(r)$). But since $U \cap C' = \emptyset$ for all predecessor components of C , we know $b \in C$ and therefore $b \in (U \cap C)$, which implies $\mathbf{A} \dot{\cup} \neg.(U \cap C) \not\models b$.

If b is a positive or default-negated external atom, then all input atoms a to b are either in a predecessor component C' of C or in C itself (since $h \rightarrow_e a$ for all $h \in H(r)$). We show with a similar argument as before that the truth value of each input atom a is the same under $\mathbf{A} \dot{\cup} \neg.U$ and $\mathbf{A} \dot{\cup} \neg.(U \cap C)$: if $\mathbf{A} \dot{\cup} \neg.U \models a$, then $\mathbf{T}a \in \mathbf{A}$ and $a \notin U$, hence $a \notin (U \cap C)$ and therefore $\mathbf{A} \dot{\cup} \neg.(U \cap C) \models a$. If $\mathbf{A} \dot{\cup} \neg.U \not\models a$, then either $\mathbf{F}a \in \mathbf{A}$, which immediatly implies $\mathbf{A} \dot{\cup} \neg.(U \cap C) \not\models a$, or $a \in U$. But in the latter case a must be in C because $U \cap C' = \emptyset$ for all predecessor components C' of C . Therefore $a \in (U \cap C)$ and consequently $\mathbf{A} \dot{\cup} \neg.(U \cap C) \not\models a$. Because all input atoms a have the same truth value under $\mathbf{A} \dot{\cup} \neg.U$ and $\mathbf{A} \dot{\cup} \neg.(U \cap C)$, the same holds also for the positive or default-negated external atom b itself. \square

Proof of Proposition 5.15. If $U = \emptyset$, then the result holds trivially. By definition of Π_C , we have $H(r) \cap C = \emptyset$ for all $r \in \Pi \setminus \Pi_C$. By hypothesis we have $U \subseteq C$. But then $H(r) \cap U = \emptyset$ for all $r \in \Pi \setminus \Pi_C$ and U is an unfounded set of Π wrt. \mathbf{A} . \square