

INSTITUT FÜR INFORMATIONSSYSTEME
ARBEITSBEREICH WISSENSBASIERTE SYSTEME

A MODEL BUILDING FRAMEWORK FOR
ANSWER SET PROGRAMMING WITH
EXTERNAL COMPUTATIONS

THOMAS EITER MICHAEL FINK
GIOVAMBATTISTA IANNI THOMAS KRENNWALLNER
CHRISTOPH REDL PETER SCHÜLLER

INFSYS RESEARCH REPORT 15-01
JANUARY 2015

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*

A MODEL BUILDING FRAMEWORK FOR ANSWER SET
PROGRAMMING WITH EXTERNAL COMPUTATIONS

Thomas Eiter¹ Michael Fink¹ Giovambattista Ianni²
Thomas Krennwallner¹ Christoph Redl¹ Peter Schüller³

Abstract. As software systems are getting increasingly connected, there is a need for equipping nonmonotonic logic programs with access to external sources that are possibly remote and may contain information in heterogeneous formats. To cater for this need, HEX programs were designed as a generalization of answer set programs with an API style interface that allows to access arbitrary external sources, providing great flexibility. Efficient evaluation of such programs however is challenging, and it requires to interleave external computation and model building; to decide when to switch between these tasks is difficult, and existing approaches have limited scalability in many real-world application scenarios. We present a new approach for the evaluation of logic programs with external source access, which is based on a configurable framework for dividing the non-ground program into possibly overlapping smaller parts called evaluation units. The latter will be processed by interleaving external evaluation and model building using an evaluation graph and a model graph, respectively, and by combining intermediate results. Experiments with our prototype implementation show a significant improvement compared to previous approaches. While designed for HEX-programs, the new evaluation approach may be deployed to related rule-based formalisms as well.

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

²Università della Calabria, 87036 Rende (CS), Italy; email: ianni@mat.unical.it.

³Faculty of Engineering, Marmara University, 34722 Istanbul, Turkey; email: peter.schuller@marmara.edu.tr

Acknowledgements: This research has been supported by the Austrian Science Fund (FWF) project P24090. Preliminary results of this work have been presented at LPNMR 2011 (12).

Copyright © 2015 by the authors

Contents

1	Introduction	4
2	Language Overview	5
2.1	HEX Syntax	5
2.2	HEX Semantics	7
2.3	Using HEX-Programs for Knowledge Representation and Reasoning	9
2.3.1	Computation Outsourcing	9
2.3.2	Knowledge Outsourcing	10
2.3.3	Combinations	10
3	Extensional Semantics and Atom Dependencies	10
3.1	Restriction to Extensional Semantics for HEX External Atoms	10
3.2	Atom Dependencies	11
3.3	Safety Restrictions	13
4	Rule Dependencies and Generalized Rule Splitting Theorem	15
4.1	Rule Dependencies	15
4.2	Splitting Sets and Theorems	16
5	Decomposition and Evaluation Techniques	18
5.1	Evaluation Graph	18
5.1.1	Evaluation Graph Splitting	20
5.1.2	First Ancestor Intersection Units	21
5.2	Interpretation Graph	22
5.2.1	Join	24
5.3	Answer Set Graph	25
5.3.1	Complete Answer Set Graphs	26
5.4	Answer Set Building	27
5.4.1	Model Streaming	29
6	Implementation	29
6.1	System Architecture	29
6.2	Heuristics	31
6.3	Experimental Results	32
6.3.1	Multi-Context Systems (MCS)	33
6.3.2	Reviewer Selection (RS)	34
6.3.3	Summary	35
7	Related Work and Discussion	35
7.1	Related Work	35
7.1.1	External Sources	35
7.1.2	Rule Dependencies	36
7.1.3	Modularity	36
7.1.4	Splitting Theorems	36

INFSYS RR 15-01	3
7.2 Possible Optimizations	37
8 Conclusion	38
8.1 Outlook	38
A Proofs	38
B Example Run of Algorithm 2	43
C On Demand Model Streaming Algorithm	44
D Overview of Liberal Domain-Expansion Safety	46

1 Introduction

Motivated by a need for knowledge bases to access external sources, extensions of declarative KR formalisms have been conceived that provide this capability, which is often realized via an API-style interface. In particular, HEX programs (19) extend nonmonotonic logic programs under the stable model semantics with the possibility to bidirectionally access external sources of knowledge and/or computation. E.g., a rule

$$pointsTo(X, Y) \leftarrow \&hasHyperlink[X](Y), url(X)$$

might be used for obtaining pairs of URLs (X, Y) , where X actually links Y on the Web, and $\&hasHyperlink$ is an *external predicate* construct. Besides constants (i.e., values) as above, also relational knowledge (predicate extensions) can flow from external sources to the logic program and vice versa, and recursion involving external predicates is allowed under safety conditions. This facilitates a variety of applications that require logic programs to interact with external environments, such as querying RDF sources using SPARQL (43), default rules on ontologies (29; 10), complaint management in e-government (55), material culture analysis (37), user interface adaptation (54), multi-context reasoning (5), or robotics and planning (50; 28), to mention a few.

Despite the absence of function symbols, an unrestricted use of external atoms leads to undecidability, as new constants may be introduced from the sources; in iteration, this can lead to an infinite Herbrand universe for the program. However, even under suitable restrictions like liberal domain-expansion safety (15) that avoid this problem, the efficient evaluation of HEX-programs is challenging, due to aspects like non-monotonic atoms and recursive access (e.g., in transitive closure computations).

Advanced in this regard was the work by 14), which fostered an evaluation approach using a traditional LP system. Roughly, the values of ground external atoms are guessed, model candidates are computed as answer sets of a rewritten program, and then those discarded which violate the guess. Compared to previous approaches such as the one by 20), it further exploits conflict-driven techniques which were extended to external sources. A generalized notion of Splitting Set (34) was introduced by 20) for non-ground HEX-programs, which were then split into subprograms with and without external access, where the former are as large and the latter as small as possible. The subprograms are evaluated with various specific techniques, depending on their structure (20; 48). However, for real-world applications this approach has severe scalability limitations, as the number of ground external atoms may be large, and their combination causes a huge number of model candidates and memory outage without any answer set output.

To remedy this problem, we reconsider model computation and make several contributions, which are summarized as follows.

- We present a modularity property of HEX-programs based on a novel generalization of the Global Splitting Theorem (20), which lifted the Splitting Set Theorem (34) to HEX-programs. In contrast to previous results, the new result is formulated on a *rule splitting set* comprising rules that may be non-ground, moreover it is based on rule dependencies rather than atom dependencies. This theorem allows for defining answer sets of the overall program in terms of the answer sets of program components that may be non-ground.
- Moreover, we present a generalized version of the new splitting theorem which allows for sharing constraints across the split; this helps to prune irrelevant partial models and candidates earlier than in previous approaches. As a consequence — and different from other decomposition approaches— subprograms for evaluation may overlap and also be non-maximal resp. non-minimal.
- Based on this theorem, we present an evaluation framework that allows for flexible evaluation of HEX-programs. It consists of an *evaluation graph* and a *model graph*; the former captures a modular decompo-

sition and partial evaluation order of the program resp. its rules, while the latter comprises for each node collections of sets of input models (which need to be combined) and output models to be passed on between components. This structure allows us to realize customized divide-and-conquer evaluation strategies. As the method works on non-ground programs, introducing new values by external calculations is feasible, as well as applying optimization based on domain splitting (13).

- A generic prototype of the evaluation framework has been implemented which can be instantiated with different solvers for Answer Set Programming (ASP) (in our suite, with `dlv` and `clasp`). It also features *model streaming*, i.e., enumeration of the models one by one. In combination with early model pruning, this can considerably reduce memory consumption and avoid termination without solution output in a larger number of settings.

Applying it to ordinary programs (without external functions) allows us to do parallel solving with a solver software that does not have parallel computing capabilities itself ('parallelize from outside').

The paper is organized as follows. In Section 2 we present the HEX-language and consider an example to demonstrate it in an intuitive way; we will use it as a running example throughout the paper. In Section 3 we then introduce necessary restrictions and preliminary concepts that form dependency-based program evaluation. After that, we develop in Section 4 our generalized splitting theorem, which is applied in Section 5 to build a new decomposition framework. Some details about the implementation and experimental results are given in Section 6. After a discussion including related work in Section 7, the paper concludes in Section 8. The proofs of all technical results have been moved to A.

2 Language Overview

In this section, we introduce the syntax and semantics of HEX-programs as far as this is necessary to explain use cases and basic modeling in the language.

2.1 HEX Syntax

Let \mathcal{C} , \mathcal{X} , and \mathcal{G} be mutually disjoint sets whose elements are called *constant names*, *variable names*, and *external predicate names*, respectively. Unless explicitly specified, elements from \mathcal{X} (resp., \mathcal{C}) are denoted with first letter in upper case (resp., lower case), while elements from \mathcal{G} are prefixed with '&'. Note that constant names serve both as individual and predicate names.

Elements from $\mathcal{C} \cup \mathcal{X}$ are called *terms*. An *atom* is a tuple (Y_0, Y_1, \dots, Y_n) , where Y_0, \dots, Y_n are terms; $n \geq 0$ is the *arity* of the atom. Intuitively, Y_0 is the predicate name, and we thus also use the more familiar notation $Y_0(Y_1, \dots, Y_n)$. The atom is *ordinary* (resp. *higher-order*), if Y_0 is a constant (resp. a variable). An atom is *ground*, if all its terms are constants. Using an auxiliary predicate aux_n for each arity n , we can easily eliminate higher-order atoms by rewriting them to ordinary atoms $aux_n(Y_0, \dots, Y_n)$. We therefore assume in the rest of this article that programs have no higher-order atoms.

An *external atom* is of the form

$$\&g[Y_1, \dots, Y_n](X_1, \dots, X_m), \tag{1}$$

where Y_1, \dots, Y_n and X_1, \dots, X_m are two lists of terms (called *input* and *output* lists, respectively), and $\&g \in \mathcal{G}$ is an external predicate name. We assume that $\&g$ has fixed lengths $in(\&g) = n$ and $out(\&g) = m$ for input and output lists, respectively.

$$P_{swim}^{EDB} = \left\{ \begin{array}{l} location(in, margB), location(in, amalB), \\ location(out, gansD), location(out, altD) \end{array} \right\}$$

$$P_{swim}^{IDB} = \left\{ \begin{array}{l} r_1: swim(in) \vee swim(out) \leftarrow . \\ r_2: need(inout, C) \leftarrow \&rq[swim](C). \\ r_3: goto(X) \vee ngoto(X) \leftarrow swim(P), location(P, X). \\ r_4: go \leftarrow goto(X). \\ r_5: need(loc, C) \leftarrow \&rq[goto](C). \\ c_6: \leftarrow goto(X), goto(Y), X \neq Y. \\ c_7: \leftarrow not\ go. \\ c_8: \leftarrow need(X, money). \end{array} \right\}$$

Figure 1: Program $P_{swim} = P_{swim}^{EDB} \cup P_{swim}^{IDB}$ to decide swimming location

Intuitively, an external atom provides a way for deciding the truth value of an output tuple depending on the input tuple and a given interpretation.

Example 1 (a, b, c) , $a(b, c)$, $node(X)$, and $D(a, b)$ are atoms; the first three are ordinary, where the second atom is a syntactic variant of the first, while the last atom is higher-order.

The external atom $\&reach[edge, a](X)$ may be devised for computing the nodes which are reachable in the graph $edge$ from the node a . We have for the input arity $in(\&reach) = 2$ and for the output arity $out(\&reach) = 1$. Intuitively, $\&reach[edge, a](X)$ will be true for all ground substitutions $X \mapsto b$ such that b is a node in the graph given by $edge$, and there is a path from a to b in that graph. \square

Definition 1 (rules and HEX programs) A rule r is of the form

$$\alpha_1 \vee \dots \vee \alpha_k \leftarrow \beta_1, \dots, \beta_n, not\ \beta_{n+1}, \dots, not\ \beta_m, \quad m, k \geq 0, \quad (2)$$

where all α_i are atoms and all β_j are either atoms or external atoms. We let $H(r) = \{\alpha_1, \dots, \alpha_k\}$ and $B(r) = B^+(r) \cup B^-(r)$, where $B^+(r) = \{\beta_1, \dots, \beta_n\}$ and $B^-(r) = \{\beta_{n+1}, \dots, \beta_m\}$. Furthermore, a (HEX) program is a finite set P of rules.

We denote by $const(P)$ the set of constant symbols occurring in a program P .

A rule r is a *constraint*, if $H(r) = \emptyset$ and $B(r) \neq \emptyset$; a *fact*, if $B(r) = \emptyset$ and $H(r) \neq \emptyset$; and *nondisjunctive*, if $|H(r)| \leq 1$. We call r *ordinary*, if it contains only ordinary atoms. We call a program P *ordinary* (resp., *nondisjunctive*), if all its rules are ordinary (resp., nondisjunctive).

Example 2 (Swimming Example) Imagine Alice wants to go for a swim in Vienna. She knows two indoor pools called Margarethenbad and Amalienbad (represented by $margB$ and $amalB$, respectively), and she knows that outdoor swimming is possible in the river Danube at two locations called Gänsehäufel and Alte Donau (denoted $gansD$ and $altD$, respectively).¹ She looks up on the Web whether she needs to pay an entrance fee, and what additional equipment she will need. Finally she has the constraint that she does not want to pay for swimming.

The HEX program $P_{swim} = P_{swim}^{EDB} \cup P_{swim}^{IDB}$ shown in Figure 1 represents Alice's reasoning problem. The extensional part P_{swim}^{EDB} contains a set of facts about possible swimming locations (where in and out are short for indoor and outdoor, respectively). The intensional part P_{swim}^{IDB} incorporates the web research of Alice in an external computation, i.e., using an external atom of the form $\&rq\langle choice \rangle\langle resource \rangle$.

¹To keep the example simple, we assume Alice knows no other possibilities to go swimming in Vienna.

Assume Alice finds out that indoor pools in general have an admission fee, and that one also has to pay at Gänsehäufel, but not at Alte Donau. Furthermore Alice reads some reviews about swimming locations and finds out that she will need her Yoga mat for Alte Donau because the ground is so hard, and she will need goggles for Amalienbad because there is so much chlorine in the water.

We next explain the intuition behind the rules in P_{swim} : r_1 chooses indoor vs. outdoor swimming locations, and r_2 collects requirements that are caused by this choice. Rule r_3 chooses one of the indoor vs. outdoor locations, depending on the choice in r_1 , and r_5 collects requirements caused by this choice. By r_4 and c_7 we ensure that some location is chosen, and by c_6 that only a single location is chosen. Finally c_8 rules out all choices that require money.

The external predicate $\&rq$ has $in(\&rq) = out(\&rq) = 1$; intuitively $\&rq[\alpha](\beta)$ is true if a resource β is required when swimming in a place in the extension of predicate α . For example, $\&rq[swim](money)$ is true if $swim(in)$ is true, because indoor swimming pool charge money for swimming. Note that this only gives an intuitive account of the semantics of $\&rq$ which will formally be defined in the following examples. \square

2.2 HEX Semantics

The semantics of HEX-programs (20; 48) generalizes the answer-set semantics (27). Let P be a HEX-program. Then the *Herbrand base* of P , denoted HB_P , is the set of all possible ground versions of atoms and external atoms occurring in P obtained by replacing variables with constants from \mathcal{C} . The grounding of a rule r , $grnd(r)$, is defined accordingly, and the grounding of P is given by $grnd(P) = \bigcup_{r \in P} grnd(r)$. Unless specified otherwise, \mathcal{X} and \mathcal{G} are implicitly given by P . Different from the ‘usual’ ASP setting, the set of constants \mathcal{C} used for grounding a program is only partially given by the program itself; in HEX, external computations may introduce new constants that are relevant for semantics of the program.

Example 3 (ctd) In P_{swim} the external atom $\&rq$ can introduce constants *yogamat* and *goggles* which are not contained in P_{swim} , but they are relevant for computing answer sets of P_{swim} . \square

An *interpretation relative to P* is any subset $I \subseteq HB_P$ containing only atoms. We say that I is a *model* of atom $a \in HB_P$, denoted $I \models a$, if $a \in I$.

With every external predicate name $\&g \in \mathcal{G}$, we associate an $(n+m+1)$ -ary Boolean function $f_{\&g}$ assigning each tuple $(I, y_1 \dots, y_n, x_1, \dots, x_m)$ either 0 or 1, where $n = in(\&g)$, $m = out(\&g)$, $I \subseteq HB_P$, and $x_i, y_j \in \mathcal{C}$. We say that $I \subseteq HB_P$ is a *model* of a ground external atom $a = \&g[y_1, \dots, y_n](x_1, \dots, x_m)$, denoted $I \models a$, if $f_{\&g}(I, y_1 \dots, y_n, x_1, \dots, x_m) = 1$.

Note that this definition of external atom semantics is very general; indeed an external atom may depend on every part of the interpretation. Therefore we will later (Section 3.1) formally restrict external computations such that they depend only on the extension of those predicates in I which are given in the input list. All examples and encodings in this work obey this restriction.

Example 4 (ctd.) The external predicate $\&rq$ in P_{swim} represents Alice’s knowledge about swimming locations as follows: for any interpretation I and some predicate (i.e., constant) α ,

$$\begin{aligned} \&rq[\alpha](money) &\text{ iff } f_{\&rq}(I, \alpha, money) = 1 &\text{ iff } \alpha(in) \in I \text{ or } \alpha(gansD) \in I, \\ \&rq[\alpha](yogamat) &\text{ iff } f_{\&rq}(I, \alpha, yogamat) = 1 &\text{ iff } \alpha(altD) \in I, \text{ and} \\ \&rq[\alpha](goggles) &\text{ iff } f_{\&rq}(I, \alpha, goggles) = 1 &\text{ iff } \alpha(amalB) \in I. \end{aligned}$$

Due to this definition of $f_{\&rq}$, it holds, e.g., that $\{swim(in)\} \models \&rq[swim](money)$. This matches the intuition about $\&rq$ indicated in the previous example. \square

Let r be a ground rule. Then we say that

- (i) I satisfies the head of r , denoted $I \models H(r)$, if $I \models a$ for some $a \in H(r)$;
- (ii) I satisfies the body of r ($I \models B(r)$), if $I \models a$ for all $a \in B^+(r)$ and $I \not\models a$ for all $a \in B^-(r)$; and
- (iii) I satisfies r ($I \models r$), if $I \models H(r)$ whenever $I \models B(r)$.

We say that I is a *model* of a HEX-program P , denoted $I \models P$, if $I \models r$ for all $r \in \text{grnd}(P)$. We call P *satisfiable*, if it has some model.

Definition 2 (answer set) *Given a HEX-program P , the FLP-reduct of P with respect to $I \subseteq \text{HB}_P$, denoted fP^I , is the set of all $r \in \text{grnd}(P)$ such that $I \models B(r)$. Then $I \subseteq \text{HB}_P$ is an answer set of P if, I is a minimal model of fP^I .*

Example 5 (ctd.) *The HEX program P_{swim} with external semantics as given in the previous example has a single answer set*

$$I = \{\text{swim}(\text{out}), \text{goto}(\text{altD}), \text{ngoto}(\text{gansD}), \text{go}, \text{need}(\text{loc}, \text{yogamat})\}.$$

(Here, and in following examples, we omit $P_{\text{swim}}^{\text{EDB}}$ from all interpretations and answer sets.) Under I , the external atom $\&\text{rq}[\text{goto}](\text{yogamat})$ is true all others ($\&\text{rq}[\text{swim}](\text{money})$, $\&\text{rq}[\text{goto}](\text{money})$, $\&\text{rq}[\text{swim}](\text{yogamat})$, ...) are false. Intuitively, answer set I tells Alice to take her Yoga mat and go for a swim to Alte Donau. \square

HEX programs (19) are a conservative extension of disjunctive (resp., normal) logic programs under the answer set semantics: answer sets of *ordinary nondisjunctive* HEX programs coincide with stable models of logic programs as proposed by (26), and answer sets of *ordinary* HEX programs coincide with stable models of disjunctive logic programs (44; 27).

The FLP-reduct as used in the HEX-semantics is equivalent to the GL-reduct, which removes the default-negated part from the remaining rules and is used for ordinary ASP programs, but the former is superior for programs with aggregates as it eliminates unintuitive answer sets. To this end, consider the following example.

Example 6 *Let P be the HEX-program*

$$\begin{aligned} p(a) &\leftarrow \text{not } \&\text{not}[p](a) \\ f &\leftarrow \text{not } p(a), \text{not } f \end{aligned}$$

where $f_{\&\text{not}}(I, p, a) = 1$ if $p(a) \notin I$ and $f_{\&\text{not}}(I, p, a) = 0$ otherwise.

The program has the answer set candidates $I_1 = \{p(a)\}$, $I_2 = \{p(a), f\}$, $I_3 = \emptyset$ and $I_4 = \{f\}$. Under the GL-reduct, we have $P^{I_1} = P^{I_2} = \{p(a) \leftarrow\}$, $P^{I_3} = \{f \leftarrow\}$ and $P^{I_4} = \emptyset$. As I_1 is a minimal model of P^{I_1} , it is a GL-answer set of P ; no other candidate is a GL-answer set. However, it is not intuitive that I_1 is an answer set as $p(a)$ supports itself. Using the FLP-reduct, we get $fP^{I_1} = \{p(a) \leftarrow \text{not } \&\text{not}[p](a)\}$. But now I_1 is not a minimal model of fP^{I_1} , as I_3 is also a model of fP^{I_1} and $I_3 \subsetneq I_1$. Similarly, one can check that I is not a minimal model of fP^I for each other candidate I ; thus under the FLP-reduct, every interpretation fails to be an answer set.

In the previous example, all answer sets of a HEX program P under the FLP-reduct are in fact minimal models of P ; this is not a coincidence but holds in general. For a study of properties of hex-programs, we refer to (19; 48; 53) and (51), where also variants and refinements of the FLP-semantics are considered, as well as the particular instance called description logic programs (see Section 2.3).

2.3 Using HEX-Programs for Knowledge Representation and Reasoning

While ASP is well-suited for many problems in artificial intelligence and was successfully applied to a range of applications (cf. e.g. 38)), modern trends computing, for instance in distributed systems and the World Wide Web, require accessing other sources of computation as well. HEX-programs cater for this need by its external atoms which provide a bidirectional interface between the logic program and other sources.

One can roughly distinguish between two main usages of external sources, which we will call *computation outsourcing*, *knowledge outsourcing*, and combinations thereof. However, we emphasize that this distinction concerns the usage in an application but both are based on the same syntactic and semantic language constructs. For each of these groups we will describe some typical use cases which serve as usage patterns for external atoms when writing HEX-programs.

2.3.1 Computation Outsourcing

Computation outsourcing means to send the definition of a subproblem to an external source and retrieve its result. The input to the external source uses predicate extensions and constants to define the problem at hand and the output terms are used to retrieve the result, which can in simple cases also be a boolean decision.

On-demand Constraints A special case of the latter case are on-demand constraints of type

$$\leftarrow \&forbidden[p_1, \dots, p_n]()$$

which eliminate certain extensions of predicates p_1, \dots, p_n . Note that the external evaluation of such a constraint can also return reasons for conflicts to the reasoner in order to restrict the search space and avoid reconstruction of the same conflict (14). This is similar to the CEGAR approach in model checking (9).

Computations which cannot (easily) be Expressed by Rules Outsourcing computations also allows for including algorithms which cannot easily or efficiently be expressed in the logic program, e.g., because they involve floating-point numbers. As a concrete example, an artificial intelligence agent for the skills and tactics game *AngryBirds* needs to perform physics simulations (8). As this requires floating point computations which can practically not be done by rules as this would either come at the costs of very limited precision or a blow-up of the grounding, HEX-programs with access to an external source for physics simulations are used.

Complexity Lifting External atoms can realize computations with a complexity higher than the complexity of ordinary ASP programs. The external atom serves than as an ‘oracle’ for deciding subprograms. While for the purpose of complexity analysis of the formalism it is often assumed that external atoms can be evaluated in polynomial time (21)², as long as external sources are decidable there is no practical reason for limiting their complexity (but of course a computation with greater complexity than polynomial time lifts the complexity results of the overall formalism as well). In fact, external sources can be other ASP- or HEX-programs. This allows for encoding other formalisms of higher complexity in HEX-programs, e.g. *abstract argumentation frameworks* (11).

²Under this assumption, deciding the existence of an answer set of a propositional HEX-program is Σ_2^P -complete.

2.3.2 Knowledge Outsourcing

In contrast, knowledge outsourcing refers to external sources which store information which needs to be imported, while reasoning itself is done in the logic program.

A typical example can be found in Web resources which provide information for import, e.g. *RDF triple stores* (32) or *geographic data* (37). More advanced use cases are *multi-context systems*, which are systems of knowledge-bases (*contexts*) that are abstracted to acceptable belief sets (roughly speaking, sets of atoms) and interlinked by *bridge rules* that range across knowledge bases (5); access to individual contexts has been provided through external atoms (4). Also sensor data, as often used when planning and executing actions in an environment, is a form of knowledge outsourcing (cf. ACTHEX (3)).

2.3.3 Combinations

It is also possible to combine the outsourcing of computations and of knowledge. A typical example are logic programs with access to description logic knowledge bases (DL KBs), called *DL-programs* (18). A DL KB does not only store information, but also provides a reasoning mechanism. This allows the logic program for formalizing queries which initiate external computations based on external knowledge and importing the results.

3 Extensional Semantics and Atom Dependencies

We now introduce additional important notions related to HEX-programs. Some of the following concepts are needed to make the formalism decidable, others prepare the basic evaluation techniques presented in later sections.

3.1 Restriction to Extensional Semantics for HEX External Atoms

To make HEX programs computable in practice, it is useful to restrict external atoms, such that their semantics depends only on extensions of predicates given in the input tuple (20). This restriction is relevant for all subsequent considerations.

Syntax Each $\&g$ is associated with an input type signature t_1, \dots, t_n such that every t_i is the type of input Y_i at position i in the input list of $\&g$. A *type* is either **const** or a non-negative integer.

Consider $\&g$, its type signature t_1, \dots, t_n , and a ground external atom $\&g[y_1, \dots, y_n](x_1, \dots, x_m)$. Then, in this setting, the signature of $\&g$ enforces certain constraints on $f_{\&g}(I, y_1, \dots, y_n, x_1, \dots, x_m)$ such that its truth value depends only on

- (a) the constant value of y_i whenever $t_i = \mathbf{const}$, and
- (b) the extension of predicate y_i , of arity t_i , in I whenever $t_i \in \mathbb{N}$.

Example 7 (ctd) Continuing Example 1, for $\&reach[edge, a](x)$, we have $t_1 = 2$ and $t_2 = \mathbf{const}$. Therefore the truth value of $\&reach[edge, a](x)$ depends on the extension of binary predicate *edge*, on the constant *a*, and on *x*.

Continuing Example 4, the external predicate $\&rq$ has $t_1 = 1$, therefore the truth value of $\&rq[swim](x)$ for various *x* wrt. an interpretation *I* depends on the extension of the unary predicate *swim* in the input list.

□

Note that the truth value of an external atom with only constant input terms, i.e., $t_i = \mathbf{const}$, $1 \leq i \leq n$, is independent of I .

Semantic constraints enforced by signatures are formalized next.

Semantics Let a be a type, I be an interpretation and $p \in \mathcal{C}$. The *projection function* $\Pi_a(I, p)$ is the binary function such that $\Pi_{\mathbf{const}}(I, p) = p$ for $a = \mathbf{const}$, and $\Pi_a(I, p) = \{(x_1, \dots, x_a) \mid p(x_1, \dots, x_a) \in I\}$ for $a \in \mathbb{N}$. Recall that atoms $p(x_1, \dots, x_a)$ are tuples (p, x_1, \dots, x_a) ; thus $D_a := \mathcal{C}^{a+1}$, i.e., the $a+1$ -fold cartesian product of \mathcal{C} , contains all syntactically possible atoms with a arguments. Furthermore, we let $D_{\mathbf{const}} := \mathcal{C}$.

Definition 3 (extensional evaluation function) Let $\&g$ be an external predicate with oracle function $f_{\&g}$, $in(\&g) = n$, $out(\&g) = m$, and type signature t_1, \dots, t_n . Then the extensional evaluation function $F_{\&g} : D_{t_1} \times \dots \times D_{t_n} \rightarrow 2^{\mathcal{C}^m}$ of $\&g$ is defined such that for every $\mathbf{a} = (a_1, \dots, a_m)$

$$\mathbf{a} \in F_{\&g}(\Pi_{t_1}(I, p_1), \dots, \Pi_{t_n}(I, p_n)) \text{ iff } f_{\&g}(I, p_1, \dots, p_n, a_1, \dots, a_m) = 1.$$

Note that $F_{\&g}$ makes the possibility of new constants in external atoms more explicit: tuples returned by $F_{\&g}$ may contain constants that are not contained in P . Furthermore, $F_{\&g}$ is well-defined only under the assertion at the beginning of this section.

Example 8 (ctd) For I from Example 5, we have $\Pi_1(I, swim) = \{(swim, out)\}$ and $\Pi_1(I, goto) = \{(goto, altD)\}$. The extensional evaluation function of $\&rq$ is

$$F_{\&rq}(U) = \{(money) \mid (X, in) \in U \text{ or } (X, gansD) \in U\} \cup \{(yogamat) \mid (X, altD) \in U\} \cup \{(goggles) \mid (X, amalB) \in U\}$$

Observe that none of the constants *yogamat* and *goggles* occurs in P (we have that $const(P) = \{swim, goto, ngoto, need, go, inout, loc, in, out, amalB, gansD, altD, margB, money, location\}$). they are introduced by the external atom semantics. Note that *(money)* is a unary tuple, as $\&rq$ has a unary output list. \square

3.2 Atom Dependencies

To account for dependencies between heads and bodies of rules is a common approach for realizing semantics of ordinary logic programs, as done, e.g., by means of the notions of *stratification* and its refinements like *local stratification* (45) or *modular stratification* (47), or by *splitting sets* (34). In HEX programs, head-body dependencies are not the only possible source of predicate interaction. Therefore new types of (non-ground) dependencies were considered by (20) and (48). In the following we recall these definitions but slightly reformulate and extend them, to prepare for the following sections where we lift atom dependencies to rule dependencies.

In contrast to the traditional notion of dependency, which in essence hinges on propositional programs, we must consider non-ground atoms; such atoms a and b clearly depend on each other if they unify, which we denote by $a \sim b$.

For analyzing program properties it is relevant whether a dependency is positive or negative. Whether the value of an external atom a depends on the presence of an atom b in an interpretation I depends in turn on the oracle function $f_{\&g}$ that is associated with the external predicate $\&g$ of a . Depending on other atoms in I , in some cases the *presence* of b might make a true, in some cases its *absence*. Therefore we

will not speak of *positive* and *negative* dependencies, as by 12), but more adequately of *monotonic* and *nonmonotonic* dependencies, respectively.³

Definition 4 An external predicate $\&g$ is *monotonic*, if for all interpretations I, I' such that $I \subseteq I'$ and all tuples of constants \mathbf{X} , $f_{\&g}(I, \mathbf{X}) = 1$ implies $f_{\&g}(I', \mathbf{X}) = 1$; otherwise $\&g$ is *nonmonotonic*. Furthermore, a ground external atom a is *monotonic*, if for all interpretations I, I' such that $I \subseteq I'$ we have $I \models a$ implies $I' \models a$; a non-ground external atom is *monotonic*, if each of its ground instances is *monotonic*.

Clearly, each external atom that involves a monotonic external predicates is monotonic, but not vice versa; thus monotonicity of external atoms is more fine-grained. In the sequel, we confine for simplicity to monotonic external predicates, which were underlying the original dependency definitions; the extension to arbitrary monotonic external atoms is straightforward.

Example 9 (ctd) Consider $F_{\&rq}(U)$ in Example 8: adding tuples to U cannot remove tuples from $F_{\&rq}(U)$, therefore $\&rq$ is a monotonic external predicate. \square

Next we define relations for dependencies from external atoms to other atoms.

Definition 5 (External Atom Dependencies) Let P be a HEX program, let $a = \&g[X_1, \dots, X_n](\mathbf{Y})$ in P be an external atom with the type signature t_1, \dots, t_n and let $b = p(\mathbf{Z})$ be an atom in the head of a rule in P . Then a depends external monotonically (resp., nonmonotonically) on b , denoted $a \rightarrow_m^e b$ (resp., $a \rightarrow_{nm}^e b$), if $\&g$ is monotonic (resp., nonmonotonic), $t_i \in \mathbb{N}$, \mathbf{Z} has arity t_i , and b is of form $p(\mathbf{Z})$ and $X_i = p$. We define that $a \rightarrow^e b$ if $a \rightarrow_m^e b$ or $a \rightarrow_{nm}^e b$.

Example 10 (ctd) In our example we have the three external dependencies $\&rq[\text{swim}](C) \rightarrow_m^e \text{swim}(\text{in})$, $\&rq[\text{swim}](C) \rightarrow_m^e \text{swim}(\text{out})$, and $\&rq[\text{goto}](C) \rightarrow_m^e \text{goto}(X)$. \square

As in ordinary ASP, atoms in HEX programs may depend on each other because of rules in the program.

Definition 6 For a HEX-program P and atoms α, β occurring in P , we say that

(a) α depends monotonically on β ($\alpha \rightarrow_m \beta$), if one of the following holds:

- (i) some rule $r \in P$ has $\alpha \in H(r)$ and $\beta \in B^+(r)$;
- (ii) there are rules $r_1, r_2 \in P$ such that $\alpha \in B(r_1)$, $\beta \in H(r_2)$, and $\alpha \sim \beta$; or
- (iii) some rule $r \in P$ has $\alpha \in H(r)$ and $\beta \in H(r)$.

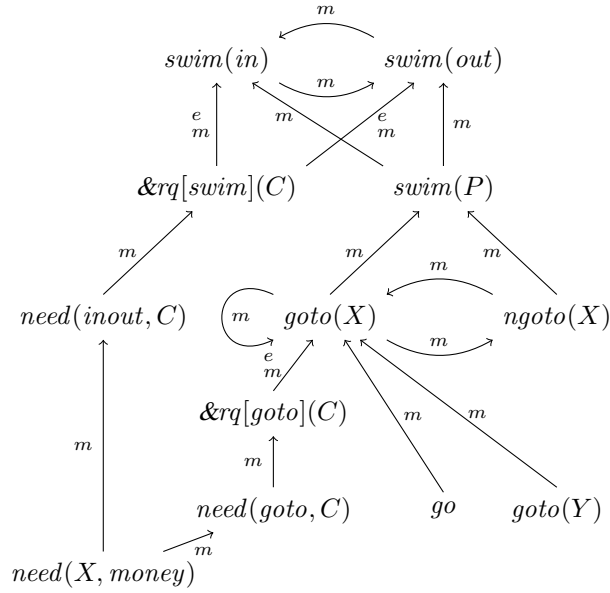
(b) α depends nonmonotonically on β ($\alpha \rightarrow_n \beta$), if there is some rule $r \in P$ such that $\alpha \in H(r)$ and $\beta \in B^-(r)$.

Note that combinations of Definitions 5 and 6 were already introduced by 48) and 13); however these works represent nonmonotonicity of external atoms in rule body dependencies and use a single ‘external dependency’ relation. In contrast, we represent nonmonotonicity of external atoms where it really happens, namely in dependencies from external atoms to ordinary atoms. We therefore obtain a simpler dependency relation between rule bodies and heads.

We say that atom α depends on atom β , denoted $\alpha \rightarrow \beta$, if either $\alpha \rightarrow_m \beta$, $\alpha \rightarrow_n \beta$, or $\alpha \rightarrow^e \beta$; that is, \rightarrow is the union of the relations \rightarrow_m , \rightarrow_n , and \rightarrow^e .

We next define the atom dependency graph.

³Note that anti-monotonicity (i.e., a larger input of an external atom can only make the external atom false, but never true) could be a third useful distinction that was exploited in (14). We here only distinguish monotonic from nonmonotonic external atoms and classify antimonotonic external atoms as nonmonotonic.

Figure 2: Atom dependency graph of running example P_{swim} .

Definition 7 For a HEX-program P , the atom dependency graph $ADG(P) = (V_A, E_A)$ of P has as vertices V_A the (non-ground) atoms occurring in non-facts of P and as edges E_A the dependency relations \rightarrow_m , \rightarrow_n , \rightarrow_m^e , and \rightarrow_{nm}^e between them in P .

Example 11 (ctd) Figure 2 shows $ADG(P_{swim})$. Note that the nonmonotonic body literal in c_7 does not show up as a nonmonotonic dependency, as c_7 has no head atoms. (The rule dependency graph in Section 4 will make this negation apparent.) \square

Next we use the dependency notions to define safety conditions on HEX programs.

3.3 Safety Restrictions

To make reasoning tasks on HEX programs decidable (or more efficiently computable), the following potential restrictions were formulated.

Rule safety This is a restriction well-known in logic programming, and it is required to ensure finite grounding of a non-ground program. A rule is safe, if all its variables are safe, and a variable is safe if it is contained in a positive body literal. Formally a rule r is safe iff variables in $H(r) \cup B^-(r)$ are a subset of variables in $B^+(r)$.

Domain-expansion safety In an ordinary logic program P , we usually assume that the set of constants \mathcal{C} is implicitly given by P . In a HEX program, external atoms may invent new constant values in their output tuples. We therefore must relax this to ‘ \mathcal{C} is countable and partially given by P ’, as shown by the following example.

Example 12 *In the Swimming Example, grounding P_{swim} with $const(P_{swim})$ is not sufficient. Further constants ‘generated’ by external atoms must be considered. For example $yogamat \notin const(P_{swim})$ and $I \models \&rq[goto](yogamat)$, hence we must ground*

$$need(loc, C) \leftarrow \&rq[goto](C)$$

with $C = yogamat$ to obtain the correct answer set. \square

Therefore grounding P with $const(P)$ can lead to incorrect results. Hence we want to obtain new constants during evaluation of external atoms, and we must use these constants to evaluate the remainder of a given HEX program. However, to ensure decidability, this process of obtaining new constants must always terminate.

Hence, we require programs to be *domain-expansion safe* (20): there must not be a cyclic dependency between rules and external atoms such that an input predicate of an external atom depends on a variable output of that same external atom, if the variable is not guarded by a domain predicate.

With HEX we need the usual notion of rule safety, i.e., a syntactic restriction which ensures that each variable in a rule only has a finite set of relevant constants for grounding. As external computations can introduce new constants in their output lists, ensuring safety in HEX is not as straightforward as in ordinary ASP.

We first recall the definition of safe variables and safe rules for HEX.

Definition 8 (Def. 5 by 20) *The safe variables of a rule r is the smallest set of variables X that occur either (i) in some ordinary atom $\beta \in B^+(r)$, or (ii) in the output list \mathbf{X} of an external atom $\&g[Y_1, \dots, Y_n](\mathbf{X})$ in $B^+(r)$ where all Y_1, \dots, Y_n are safe. A rule r is safe, if each variable in r is safe.⁴*

However, safety alone does not guarantee finite grounding of HEX programs, because an external atom might create new constants, i.e., constants not part of the program itself, in its output list (see Example 8). These constants can become part of the extension of an atom in the rule head, and by grounding and evaluation of other rules become part of the extension of a predicate which is an input to the very same external atom.

Example 13 (adapted from 48)) *The following HEX program is safe according to Definition 8 and nevertheless cannot be finitely grounded:*

$$\begin{aligned} source(\text{“http://some_url”}) &\leftarrow . \\ url(X) &\leftarrow \&rdf[source](X, \text{“rdf:subClassOf”}, C). \\ source(X) &\leftarrow url(X). \end{aligned}$$

Suppose the $\&rdf[source](S, P, O)$ atom retrieves all triples (S, P, O) from all RDF triplestores specified in the extension of $source$, and suppose that each triplestore contains a triple with a URL S that does not show up in another triplestore. As a result, all these URLs are collected in the extension of $source$ which leads to even more URLs being retrieved and a potentially infinite grounding.

However, we could change the rule with the external atom to

$$url(X) \leftarrow \&rdf[source](X, \text{“rdf:subClassOf”}, C), limit(X) \tag{3}$$

and add an appropriate set of *limit* facts. This addition of a range predicate $limit(X)$ which does not depend on the external atom output ensures a finite grounding. \square

⁴This is stated by 20) as ‘if each variable appearing in a negated atom and in any input list is safe, and variables appearing in $H(r)$ are safe’, which is equivalent.

To obtain a syntactic restriction that ensures finite grounding for HEX, so called *strong safety* has been introduced for the HEX programs (20). Intuitively, this concept requires all output variables of cyclic external atoms (using the dependency notion from Definition 7) to be bounded by ordinary body atoms of the same rule which are not part of the cycle. However, this condition is unnecessarily restrictive, and Therefore, the extensible notion of *liberal domain-expansion safety (lde-safety)* was introduced by 15), which we will use in the following. For the purpose of this article, we may omit the formal details of lde-safety (see 15) and D for an outline); it is sufficient to know that every lde-safe program has a finite grounding that has the same answer sets as the original program.

4 Rule Dependencies and Generalized Rule Splitting Theorem

In this section, we introduce a new notion of dependencies in HEX-programs, namely between non-ground *rules* in a program. Based on this notion, we then present a modularity property of HEX-programs that allows us to obtain answer sets of a program from the answer sets of its components. The property is formulated as a splitting theorem based on dependencies among rules and lifts a similar result for dependencies among atoms, viz. the Global Splitting Theorem (20), to this setting, and it generalizes and improves it. This result is exploited in a more efficient HEX-program evaluation algorithm, which we show in Section 5.

4.1 Rule Dependencies

We define rule dependencies as follows.

Definition 9 (Rule dependencies) *Let P be a program and a, b atoms occurring in distinct rules $r, s \in P$. Then r depends on s according to the following cases:*

- (i) *if $a \sim b$, $a \in B^+(r)$, and $b \in H(s)$, then $r \rightarrow_m s$;*
- (ii) *if $a \sim b$, $a \in B^-(r)$, and $b \in H(s)$, then $r \rightarrow_n s$;*
- (iii) *if $a \sim b$, $a \in H(r)$, and $b \in H(s)$, then both $r \rightarrow_m s$ and $s \rightarrow_m r$;*
- (iv) *if $a \rightarrow^e b$, $a \in B(r)$ is an external atom, and $b \in H(s)$, then*
 - $r \rightarrow_m s$ *if $a \in B^+(r)$ and $a \rightarrow_m^e b$, and*
 - $r \rightarrow_n s$ *otherwise.*

Intuitively, conditions (i) and (ii) reflect the fact that the applicability of a rule r depends on the applicability of a rule s with a head that unifies with a literal in the body of rule r ; condition (iii) exists because r and s cannot be evaluated independently if they share a common head atom (e.g., $u \vee v \leftarrow$ cannot be evaluated independently from $v \vee w \leftarrow$); and (iv) defines dependencies due to predicate inputs of external atoms.

In the sequel, we let $\rightarrow_{m,n} = \rightarrow_m \cup \rightarrow_n$ be the union of monotonic and nonmonotonic rule dependencies. We next define graphs of rule dependencies.

Definition 10 *Given a HEX-program P , the rule dependency graph $DG(P) = (V_D, E_D)$ of P is the labeled graph with vertex set $V_D = P$ and edge set $E_D = \rightarrow_{m,n}$.*

Example 14 (ctd.) *Figure 3 depicts the rule dependency graph of our running example. According to Definition 9, we have the following rule dependencies in P_{swim}^{IDB} :*

- *due to (i) we have $r_3 \rightarrow_m r_1$, $r_4 \rightarrow_m r_3$, $c_6 \rightarrow_m r_3$, $c_8 \rightarrow_m r_2$, and $c_8 \rightarrow_m r_5$;*
- *due to (ii) we have $c_7 \rightarrow_n r_4$;*

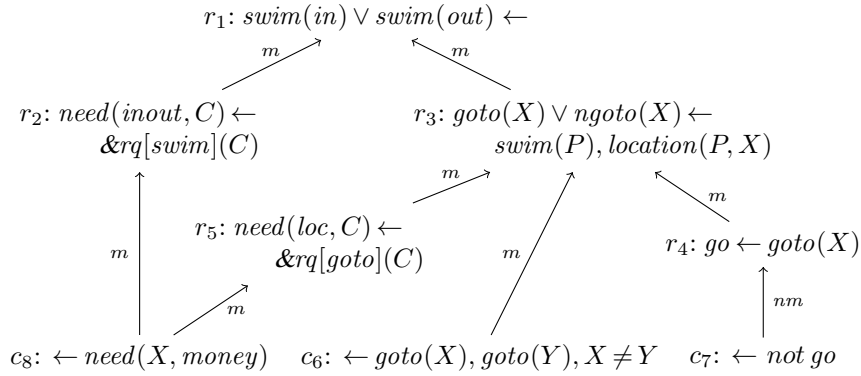


Figure 3: Rule dependency graph of running example P_{swim} .

- due to (iii) we have no dependencies; and
- due to (iv) we have $r_2 \rightarrow_m r_1$ and $r_5 \rightarrow_m r_3$.

Note that $\&rq$ is monotonic (see Example 9). \square

4.2 Splitting Sets and Theorems

Splitting sets are a notion that allows for describing how a program can be decomposed into parts and how semantics of the overall program can be obtained from semantics of these parts in a divide-and-conquer manner.

We lift the original HEX splitting theorem (20, Theorem 2) and the according definitions of global splitting set, global bottom, and global residual (20, Definitions 8 and 9) to our new definition of dependencies among rules.

A *rule splitting set* is a part of a (non-ground) program that does not depend on the rest of the program. This corresponds in a sense with global splitting sets by 20).

Definition 11 (Rule Splitting Set) A rule splitting set R for a HEX-program P is a set $R \subseteq P$ of rules such that whenever $r \in R$, $s \in P$, and $r \rightarrow_{m,n} s$, then $s \in R$ holds.

Example 15 (ctd) The following are some rule splitting sets of P_{swim} : $\{r_1\}$, $\{r_1, r_2\}$, $\{r_1, r_3\}$, $\{r_1, r_2, r_3\}$, $\{r_1, r_2, r_3, r_5, c_8\}$. The set $R = \{r_1, r_2, c_8\}$ is not a rule splitting set, because $c_8 \rightarrow_m r_5$ but $r_5 \notin R$. \square

Because of possible constraint duplication, we no longer partition the input program, and the customary notion of splitting set, bottom, and residual, is not appropriate for sharing constraints between bottom and residual. Instead, we next define a *generalized bottom* of a program, which splits a non-ground program into two parts which may share certain constraints.

Definition 12 (Generalized Bottom) Given a rule splitting set R of a HEX-program P , a generalized bottom B of P wrt. R is a set B with $R \subseteq B \subseteq P$ such that all rules in $B \setminus R$ are constraints that do not depend nonmonotonically on any rule in $P \setminus B$.

Example 16 (ctd) A rule splitting set R of P_{swim} (e.g., those given in Example 15) is also a generalized bottom of P_{swim} wrt. R . The set $\{r_1, r_2, c_8\}$ is not a rule splitting set, but it is a generalized bottom of P_{swim} wrt. the rule splitting set $\{r_1, r_2\}$, as c_8 is a constraint that depends only monotonically on rules in $P_{swim} \setminus \{r_1, r_2, c_8\}$. \square

Next, we describe how interpretations of a generalized bottom B of a program P lead to interpretations of P without re-evaluating rules in B . This is a generalization of the Splitting Set Theorem (34), of a modularity result for disjunctive logic programs (17, Lemma 5.1) and of the splitting theorem for (non-ground) HEX-programs by 20) (Global Splitting Theorem) and by 48) (Theorem 4.6.2).

Intuitively, this is a relaxation of the previous non-ground HEX splitting theorem: a constraint may be put both in the bottom and in the residual if it has no nonmonotonic dependencies to the residual. The benefit of such constraint sharing is a smaller number of answer sets of the bottom, and hence of fewer evaluations of the residual program.

Notation. For any set I of ground ordinary atoms, we denote by $facts(I)$ the corresponding set of ground facts; furthermore, for any set P of rules, we denote by $gh(P)$ the set of ground head atoms occurring in $grnd(P)$.

Theorem 1 (Splitting Theorem) *Given a HEX-program P and a rule splitting set R of P , $M \in \mathcal{AS}(P)$ iff $M \in \mathcal{AS}(P \setminus R \cup facts(X))$ with $X \in \mathcal{AS}(R)$.*

Using the definition of generalized bottom, we generalize the above theorem.

Theorem 2 (Generalized Splitting Theorem) *Let P be a HEX-program, let R be a rule splitting set of P , and let B be a generalized bottom of P wrt. R . Then*

$$M \in \mathcal{AS}(P) \text{ iff } M \in \mathcal{AS}(P \setminus R \cup facts(X)) \text{ where } X \in \mathcal{AS}(B).$$

Note that $B \setminus R$ contains shareable constraints that are used twice in the Generalized Splitting Theorem, viz. in computing X and in computing M .

The Generalized Splitting Theorem is useful for early elimination of answer sets of the bottom thanks to constraints which depend on it but also on rule heads outside the bottom. Such constraints can be shared between the bottom and the remaining program.

Example 17 (ctd) *We apply Theorems 1 and 2 to P_{swim} and compare them. Using the rule splitting set $\{r_1, r_2\}$, we can obtain $\mathcal{AS}(P_{swim})$ by first computing $\mathcal{AS}(\{r_1, r_2\}) = \{I_1, I_2\}$ where $I_1 = \{swim(in), need(inout, money)\}$, $I_2 = \{swim(out)\}$, and by then using Theorem 1: $X \in \mathcal{AS}(P_{swim})$ iff it holds that $X \in \mathcal{AS}(\{r_3, r_4, r_5, c_6, c_7, c_8\} \cup facts(I_1))$ or $X \in \mathcal{AS}(\{r_3, r_4, r_5, c_6, c_7, c_8\} \cup facts(I_2))$. Note that the computation with I_1 yields no answer set, as $need(inout, money) \in I_1$ satisfies the body of c_8 and “kills” any model candidate. In contrast, if we use the generalized bottom $\{r_1, r_2, c_8\}$, we have $\mathcal{AS}(\{r_1, r_2, c_8\}) = \{\{swim(out)\}\}$ and can use Theorem 2 to obtain $\mathcal{AS}(P_{swim})$ with only one further answer set computation: $X \in \mathcal{AS}(P_{swim})$ iff $X \in \mathcal{AS}(\{r_3, r_4, r_5, c_6, c_7, c_8\} \cup \{swim(out) \leftarrow\})$. Note that we use c_8 in both computations, i.e., c_8 is shared between the generalized bottom and the remaining computation. \square*

Armed with the results of this section, we proceed to program evaluation in the next section. A discussion of the new splitting theorems that compares them to previous related theorems and argues for their advantage is given in Section 7.1. However, we can summarize some points as follows.

- By moving from atom to rule splitting sets, no separate definition of the bottom is needed, which just becomes the (rule) splitting set.
- As regards HEX-programs, splitting is simple (and not troubled) if all atoms that are true in an answer set of the bottom also appear in the residual program. Typically, this is not the case in results from the literature.

- Finally, also the residual program itself is simpler (and easier to construct), by just dropping rules and adding facts. No rule rewriting needs to be done, and no extra facts need to be introduced in the residual program nor in the bottom.

The only (negligible) disadvantage of the new theorems is that the answer sets of the bottom and the residual program may no longer be disjoint; however, each residual answer set includes some (unique) bottom answer set.

5 Decomposition and Evaluation Techniques

We now introduce our new HEX evaluation framework, which is based on selections of sets of rules of a program that we call *evaluation units* (or briefly *units*).

The traditional HEX evaluation algorithm (20) uses a dependency graph over (non-ground) atoms, and gradually evaluates sets of rules (the ‘bottoms’ of a program) that are chosen based on this graph. In contrast our new evaluation algorithm exploits the rule-based modularity results for HEX-programs in Section 4.

While previously a constraint can only kill models once all its dependencies on rules are fulfilled, the new algorithm increases evaluation efficiency by sharing non-ground constraints, such that they may kill models earlier; this is safe if all their nonmonotonic dependencies are fulfilled. Moreover, units no longer must be maximal. Instead, we require that partial models of units, i.e., atoms in heads of their rules, do not interfere with those of other units. This allows for independence, efficient storage, and easy composition of partial models of distinct units.

5.1 Evaluation Graph

Using rule dependencies, we next define the notion of evaluation graph on evaluation units. We then relate evaluation graphs to splitting sets (34) and show how to use them to evaluate HEX-programs by evaluating units and combining the results.

We define evaluation units as follows.

Definition 13 *An evaluation unit (in short ‘unit’) is any lde-safe HEX-program.*

The formal definition of lde-safety (see D and 15)) is not crucial here, merely the property that a unit has a finite grounding with the same answer sets as the original program which can be effectively computed; lde-safe HEX-programs are the most general class of HEX-programs with this property and computational support.

An important point of the notion of evaluation graph is that rule dependencies $r \rightarrow_x s$ lead to different edges, i.e., unit dependencies, depending on the dependency type $x \in \{n, m\}$ and whether r resp. s is a constraint; constraints cannot (directly) make atoms true, hence they can be shared between units in certain cases, while sharing non-constraints could violate modularity.

Given a rule $r \in P$ and a set U of units, we denote by $U|_r = \{u \in U \mid r \in u\}$ the set of units that contain rule r .

Definition 14 (Evaluation graph) *An evaluation graph $\mathcal{E} = (U, E)$ of a program P is a directed acyclic graph whose vertices U are evaluation units and which fulfills the following properties:*

- (a) $P = \bigcup_{u \in U} u$, i.e., every rule $r \in P$ is contained in at least one unit;

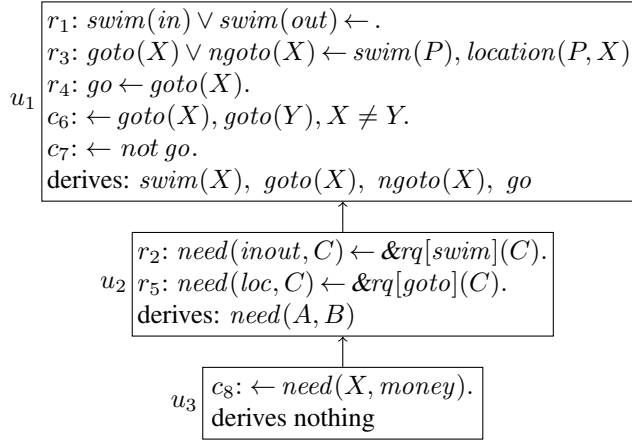


Figure 4: Evaluation graph \mathcal{E}_1 for running example HEX program P_{swim} .

- (b) every non-constraint $r \in P$ is contained in exactly one unit, i.e., $|U|_r| = 1$;
- (c) for each nonmonotonic dependency $r \rightarrow_n s$ between rules $r, s \in P$ and for all $u \in U|_r, v \in U|_s, u \neq v$, there exists an edge $(u, v) \in E$ (intuitively, nonmonotonic dependencies between rules have corresponding edges everywhere in \mathcal{E}); and
- (d) for each monotonic dependency $r \rightarrow_m s$ between rules $r, s \in P$, there exists some $u \in U|_r$ such that E contains all edges (u, v) with $v \in U|_s$ for $v \neq u$ (intuitively, for each rule r there is (at least) one unit in \mathcal{E} where all monotonic dependencies from r to other rules have corresponding outgoing edges in \mathcal{E}).

We remark that 12) and 49) defined evaluation units as *extended pre-groundable* HEX-programs; later, 46) and 15) defined *generalized evaluation units* as *lde-safe* HEX-programs, which subsume *extended pre-groundable* HEX-programs, and *generalized evaluation graphs* on top as in Definition 14. As more the grounding properties of units matter than the precise fragment, we dropped here ‘generalized’ to avoid complex terminology.

As a non-constraint can occur only in a single unit, the above definition implies that all dependencies of non-constraints have corresponding edges in \mathcal{E} , which is formally expressed in the following proposition.

Proposition 1 *Let $\mathcal{E} = (U, E)$ be an evaluation graph of a program P , and assume $r \rightarrow_{m,n} s$ is a dependency between a non-constraint $r \in P$ and a rule $s \in P$. Then $\{(u, v) \mid u \in U|_r, v \in U|_s\} \subseteq E$ holds.*

Example 18 (ctd) *Figures 4 and 5 show two possible evaluation graphs for our running example. The evaluation graph \mathcal{E}_1 contains every rule of P_{swim} in exactly one unit. In contrast, \mathcal{E}_2 contains c_8 both in u_2 and in u_4 . Condition (d) of Definition 14 is particularly interesting for these two graphs; it is fulfilled as follows. Graph \mathcal{E}_1 can be obtained by contracting rules in the rule dependency graph $DG(P_{swim})$ into units, i.e., \mathcal{E}_1 is a (graph) minor of $DG(P_{swim})$ and therefore all rule dependencies are realized as unit dependencies and Conditions (c) and (d) are satisfied. In contrast, \mathcal{E}_2 is not a minor of $DG(P_{swim})$ because dependency $c_8 \rightarrow_m r_5$ is not realized as a dependency from u_2 to u_4 . Nonetheless, all dependencies from c_8 are realized at u_4 and thus \mathcal{E}_2 conforms with condition (d), which merely requires that rule dependencies have edges corresponding to all monotonic rule dependencies at some unit of the evaluation graph. \square*

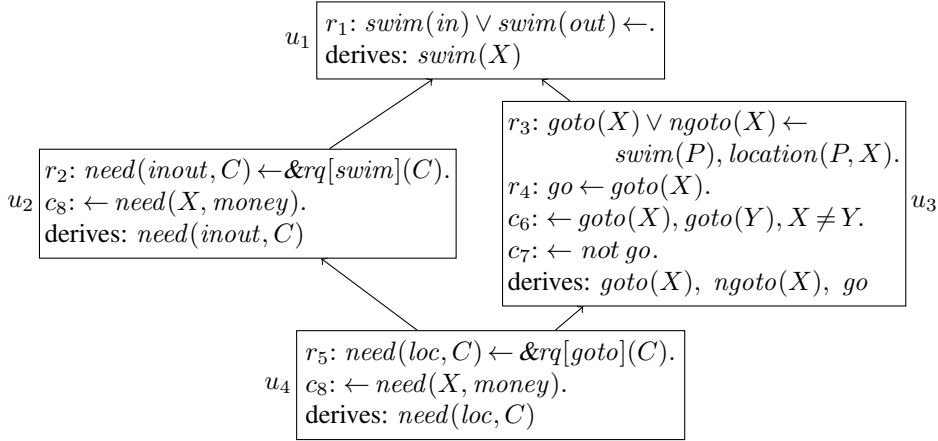


Figure 5: Evaluation graph \mathcal{E}_2 for running example HEX program P_{swim} .

Evaluation graphs have the important property that partial models of evaluation units do not intersect, i.e., evaluation units do not mutually depend on each other. This is achieved by acyclicity and because rule dependencies are covered in the graph.

In fact, due to acyclicity, mutually dependent rules of a program are contained in the same unit; thus each strongly connected component of the program's dependency graph is fully contained in a single unit. Furthermore, a unit can have in its rule heads only atoms that do not unify with atoms in the rule heads of other units, as rules which have unifiable heads mutually depend on one another. This ensures that under any grounding, the following property holds.

Proposition 2 (Disjoint unit outputs) *Let $\mathcal{E} = (U, E)$ be an evaluation graph of a program P . Then for each distinct units $u_1, u_2 \in U$, it holds that $gh(u_1) \cap gh(u_2) = \emptyset$.⁵*

We call this the property of *disjoint unit outputs*.

Example 19 (ctd) *Figures 4 and 5 show for each unit which atoms can become true due to rule heads in them, denoted as 'derived' atoms. Observe that both graphs have strictly non-intersecting atoms in rule heads of distinct units. \square*

As units of evaluation graphs can be arbitrary lde-safe programs, we clearly have the following property.

Proposition 3 *For every lde-safe HEX program P , some evaluation graph \mathcal{E} exists.*

Indeed, we can simply put P into a single unit to obtain a valid evaluation graph. Thus the HEX evaluation approach based on evaluation graphs is applicable to all domain-expansion safe HEX programs.

5.1.1 Evaluation Graph Splitting

We next show that units and their predecessors in an evaluation graph correspond to generalized bottoms. We then use this property to formulate an algorithm for unit-based, efficient evaluation of HEX-programs.

⁵See page 17 for the definition of notation $gh(P)$.




Figure 6: First Ancestor Intersection units (FAIs) in an evaluation graph.

Given an evaluation graph $\mathcal{E} = (U, E)$, we write $u < w$, if a path from u to w exists in \mathcal{E} , and $u \leq w$ if either $u < w$ or $u = w$.

For a unit $u \in U$, we denote by $preds_{\mathcal{E}}(u) = \{v \in U \mid (u, v) \in E\}$ the set of units on which u (directly) depends and by $u^{<} = \bigcup_{w \in U, u < w} w$ the set of rules in all units on which u transitively depends; furthermore, we let $u^{\leq} = u^{<} \cup u$. Note that for a leaf unit u (i.e., u has no predecessors) we have $preds_{\mathcal{E}}(u) = u^{<} = \emptyset$ and $u^{\leq} = u$.

Theorem 3 For every evaluation graph $\mathcal{E} = (U, E)$ of a HEX-program Q and unit $u \in U$, it holds that $u^{<}$ is a generalized bottom of u^{\leq} wrt. $R = \{r \in u^{<} \mid B(r) \neq \emptyset\}$.

Example 20 (ctd) In \mathcal{E}_1 , $u_2^{<} = u_1$ and $u_2^{\leq} = u_1 \cup u_2$ and $u_2^{<}$ is a generalized bottom of u_2^{\leq} wrt. $R = \{r_1, r_2, r_4\}$. In \mathcal{E}_2 , we have $u_4^{<} = u_1 \cup u_2 \cup u_3$ and $u_4^{\leq} = P_{swim}$ and $u_4^{<}$ is a generalized bottom of P_{swim} wrt. $R = \{r_1, r_2, r_3, r_4\}$. We can verify this on Definition 12: we have $P = P_{swim}$, $B = u_4^{<} = \{r_1, r_2, r_3, r_4, c_6, c_7, c_8\}$, and R as above. Then $R \subseteq B \subseteq P$, and furthermore $B \setminus R = \{c_6, c_7, c_8\}$ consists of constraints none of which depends nonmonotonically on a rule in $P \setminus B = \{r_5\}$. \square

Theorem 4 Let $\mathcal{E} = (U, E)$ be an evaluation graph of a HEX-program Q and $u \in U$. Then for every unit $u' \in preds_{\mathcal{E}}(u)$, it holds that u'^{\leq} is a generalized bottom of the subprogram $u^{<}$ wrt. the rule splitting set $R = \{r \in u'^{\leq} \mid B(r) \neq \emptyset\}$.

Example 21 (ctd) In \mathcal{E}_1 , we have $u_1 \in preds_{\mathcal{E}_1}(u_2)$; hence $u_1^{<} = u_1$ is by Theorem 4 a generalized bottom of $u_2^{<} = u_1$ wrt. $R = \{r_1, r_3, r_4\}$. Furthermore, $u_2 \in preds_{\mathcal{E}_1}(u_3)$ and hence $u_2^{<} = u_1 \cup u_2$ is a generalized bottom of $u_3^{<} = u_1 \cup u_2$ wrt. $R = \{r_1, r_2, r_3, r_4, r_5\}$. The case of \mathcal{E}_2 and u_4 is less clear. We have $u_2 \in preds_{\mathcal{E}_2}(u_4)$, thus by Theorem 4 $u_2^{<} = u_1 \cup u_2 = \{r_1, r_2, c_8\}$ is a generalized bottom of $u_4^{<} = u_1 \cup u_2 \cup u_3$ wrt. $R = \{r_1, r_2\}$. Comparing against Definition 12, we have $P = u_1 \cup u_2 \cup u_3$ and $B = u_1 \cup u_2$; thus indeed $R \subseteq B \subseteq P$ and no constraint in $B \setminus R = \{c_8\}$ depends nonmonotonically on any rule in $P \setminus B = \{r_3, r_4, c_6, c_7\}$. \square

5.1.2 First Ancestor Intersection Units

We will use the evaluation graph for model building; as syntactic dependencies reflect semantic dependencies between units, multiple paths between units need attention. Of particular importance are *first ancestor intersection units*, which are units where distinct paths starting at some unit meet first. More formally,

Definition 15 Given an evaluation graph $\mathcal{E} = (U, E)$ and units $v \neq w \in U$, we say that unit w is a first ancestor intersection unit (FAI) of v , if paths $p_1 \neq p_2$ from v to w exist in E that overlap only in v and w . By $fai(v)$ we denote the set of all FAIs of v .

Example 22 Figure 6 sketches an evaluation graph with dependencies $a \rightarrow b \rightarrow c \rightarrow e \rightarrow f$, $a \rightarrow d \rightarrow e \rightarrow g$, and $b \rightarrow d$. We have that $fai(a) = \{d, e\}$, $fai(b) = \{e\}$, and $fai(u) = \emptyset$ for each $u \in U \setminus \{a, b\}$. In particular, f and g are not FAIs of b , because all pairs of distinct paths from b to f or g overlap in more than two units. \square

Note that for tree-shaped evaluation graphs, $fai(v) = \emptyset$ for each unit v as paths between nodes in a tree are unique.

Example 23 (ctd) The evaluation graph \mathcal{E}_1 of P_{swim} is a tree (see Fig. 4), thus $fai(u) = \emptyset$ for $u \in \{u_1, u_2, u_3\}$. In contrast, the evaluation graph \mathcal{E}_2 of P_{swim} (see Fig. 5) is not a tree; we have that $fai(u_4) = \{u_1\}$ and no other unit in \mathcal{E}_2 has FAIs. \square

We can build an evaluation graph \mathcal{E} for a program P based on the dependency graph $DG(P)$. Initially, the units are set to the maximal strongly connected components of $DG(P)$, and then units are iteratively merged while preserving acyclicity and the conditions (a)-(d) of an evaluation graph; we will discuss some existing heuristics in Section 6.2, while for details we refer to 46).

5.2 Interpretation Graph

We now define the Interpretation Graph (short i-graph), which is the foundation of our model building algorithm. An i-graph is a labeled directed graph defined wrt. an evaluation graph, where each vertex is associated with a specific evaluation unit, a type (input resp. output interpretation) and an set of ground atoms.

We do not use interpretations themselves as vertices, as distinct vertices may be associated with the same interpretation; still we call vertices of the i-graph interpretations.

We first define an auxiliary concept, formulate then conditions on it, and finally define the i-graph using these conditions.

Definition 16 (Interpretation Structure) Let $\mathcal{E} = (U, E)$ be an evaluation graph for a program P . An interpretation structure \mathcal{I} for \mathcal{E} is a directed acyclic graph $\mathcal{I} = (M, F, unit, type, int)$ where $M \subseteq \mathcal{I}_{id}$ is from a countable set \mathcal{I}_{id} of identifiers, and $unit: M \rightarrow U$, $type: M \rightarrow \{1, 0\}$, and $int: M \rightarrow 2^{HB_P}$ are total node labeling functions.

The following notation will be useful. Given unit $u \in U$ in the evaluation graph associated with an i-graph \mathcal{I} , we denote by $i-ints_{\mathcal{I}}(u) = \{m \in M \mid unit(m) = u \text{ and } type(m) = 1\}$ the input (i-)interpretations, and by $o-ints_{\mathcal{I}}(u) = \{m \in M \mid unit(m) = u \text{ and } type(m) = 0\}$ the output (o-)interpretations of \mathcal{I} at unit u . For every vertex $m \in M$, we denote by

$$int^+(m) = int(m) \cup \bigcup \{int(m') \mid m' \in M \text{ and } m' \text{ is reachable from } m \text{ in } \mathcal{I}\}$$

the *expanded interpretation* of m .

Given an interpretation structure $\mathcal{I} = (M, F, unit, type, int)$ for $\mathcal{E} = (U, E)$ and a unit $u \in U$, we define the following properties:




Figure 7: Interpretation Graphs: violation of the FAI condition on the left, correct situation on the right.

(IG-I) *I-connectedness*: for every $m \in o\text{-ints}_{\mathcal{I}}(u)$, $|\{m' \mid (m, m') \in F\}| = 1$ and $m' \in i\text{-ints}_{\mathcal{I}}(u)$ is an i-interpretation at unit u ;

(IG-O) *O-connectedness*: for every $m \in i\text{-ints}_{\mathcal{I}}(u)$ and $u_i \in \text{preds}_{\mathcal{E}}(u)$, $|\{m_i \mid (m, m_i) \in F\}| = 1$ and $m_i \in o\text{-ints}_{\mathcal{I}}(u_i)$ (every m_i is an o-interpretation at u_i);

(IG-F) *FAI intersection*: let \mathcal{E}' be the subgraph of \mathcal{E} on the units reachable from u^6 and for every $m \in i\text{-ints}_{\mathcal{I}}(u)$, let \mathcal{I}' be the subgraph of \mathcal{I} reachable from m . Then \mathcal{I}' contains exactly one o-interpretation at each unit of \mathcal{E}' . (Note that both \mathcal{I} and \mathcal{E} are acyclic, hence \mathcal{I}' does not include m and \mathcal{E}' does not include u .)

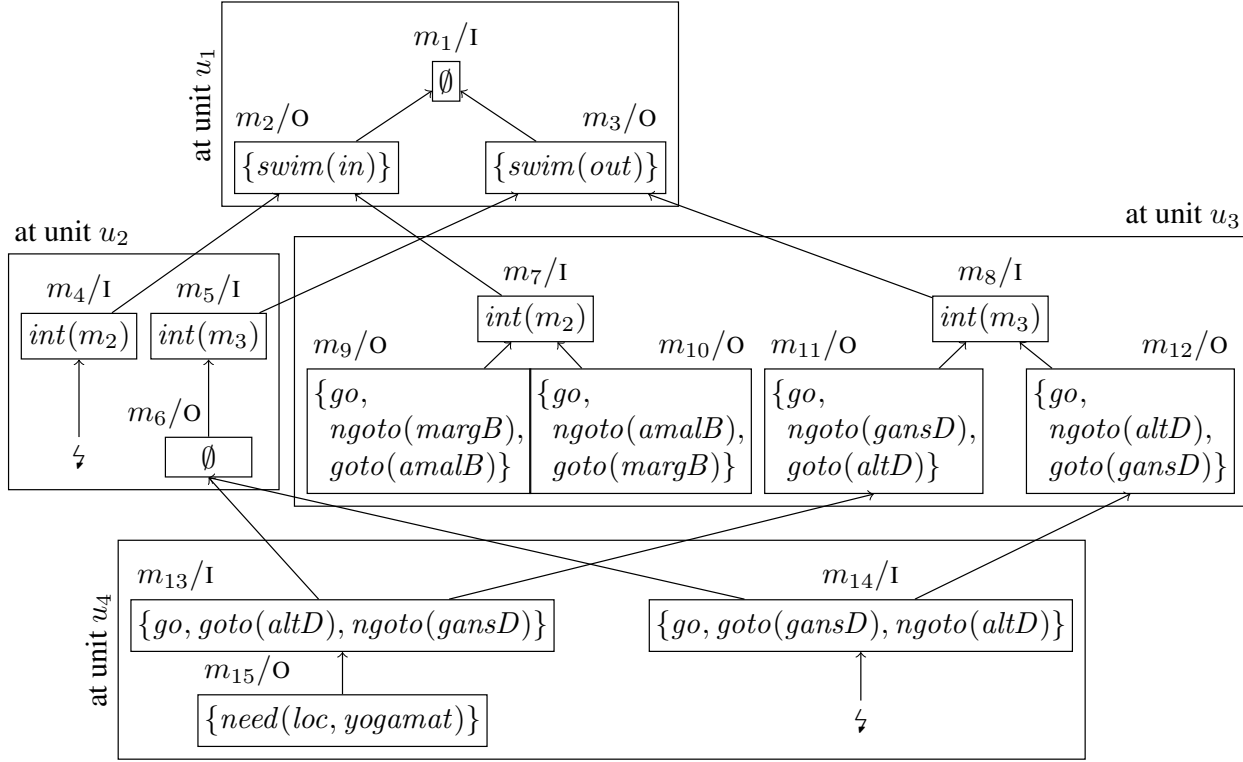
(IG-U) *Uniqueness*: for every $m_1 \neq m_2 \in M$ such that $\text{unit}(m_1) = \text{unit}(m_2) = u$, we have $\text{int}^+(m_1) \neq \text{int}^+(m_2)$ (the expanded interpretations differ).

Definition 17 (Interpretation Graph) Let $\mathcal{E} = (U, E)$ be an evaluation graph for a program P . then an interpretation graph (i-graph) for \mathcal{E} is an interpretation structure $\mathcal{I} = (M, F, \text{unit}, \text{type}, \text{int})$ that fulfills for every unit $u \in U$ the conditions (IG-I), (IG-O), (IG-F), and (IG-U).

Intuitively, the conditions make every i-graph ‘live’ on its associated evaluation graph: an i-interpretation must conform to all dependencies of the unit it belongs to, by depending on exactly one o-interpretation at that unit’s predecessor units (IG-I); moreover an o-interpretation must depend on exactly one i-interpretation at the same unit (IG-O). Furthermore, every i-interpretation depends directly or indirectly on exactly one o-interpretation at each unit it can reach in the i-graph (IG-F); this ensures that no expanded interpretation $\text{int}^+(m)$ ‘mixes’ two or more i-interpretations resp. o-interpretations from the same unit. (The effect of condition (IG-F) is visualized in Figure 7.) Finally, redundancies in an i-graph are ruled out by the uniqueness condition (IG-U).

Example 24 (ctd.) Figure 8 shows an interpretation graph \mathcal{I}_2 for \mathcal{E}_2 . (the symbol \downarrow is explained in Example 30 below). The unit label is depicted as rectangle labeled with the respective unit. The type label is indicated after interpretation names, i.e., m_1/\downarrow denotes that interpretation m_1 is an input interpretation. For \mathcal{I}_2 the set \mathcal{I}_{id} of identifiers is $\{m_1, \dots, m_{15}\}$.

⁶I.e., \mathcal{E}' is the subgraph of \mathcal{E} induced by the set of units reachable from u , including u ; in abuse of terminology, we briefly say “the subgraph (of \mathcal{E}) reachable from”

Figure 8: Interpretation graph \mathcal{I}_2 for \mathcal{E}_2

Dependencies are shown as arrows between interpretations. Observe that *I*-connectedness is fulfilled, as every *o*-interpretation depends on exactly one *i*-interpretation at the same unit. *O*-connectedness is similarly fulfilled, in particular consider *i*-interpretations of u_4 in \mathcal{I}_2 : u_4 has two predecessor units (u_2 and u_3) and every *i*-interpretation at u_4 depends on exactly one *o*-interpretation at u_2 and exactly one *o*-interpretation at u_3 . The condition on FAI intersection could only be violated by *i*-interpretations at u_4 . We can verify that from both m_{13} and m_{14} we can reach exactly one *o*-interpretation at each unit; hence the condition is fulfilled. Uniqueness is satisfied, as in both graphs no unit has two output models with the same content. \square

Note that the empty graph is an *i*-graph. This is by intent, as our model building algorithm will progress from an empty *i*-graph to one with interpretations at every unit, precisely if the program has an answer set.

5.2.1 Join

We will build *i*-graphs by adding one vertex at a time, always preserving the *i*-graph conditions. Adding an *o*-interpretation requires to add a dependency to one *i*-interpretation at the same unit. Adding an *i*-interpretation similarly requires addition of dependencies. However this is more involved because condition (IG-F) could be violated. Therefore, we next define an operation that captures all necessary conditions.

We call the combination of *o*-interpretations which yields an *i*-interpretation a ‘*join*’. Formally, the join operation ‘ \bowtie ’ is defined as follows.

Definition 18 Let $\mathcal{I} = (M, F, \text{unit}, \text{type}, \text{int})$ be an *i*-graph for an evaluation graph $\mathcal{E} = (V, E)$ of a program P . Let $u \in V$ be a unit, let $\text{preds}_{\mathcal{E}}(u) = \{u_1, \dots, u_k\}$ be the predecessor units of u , and let $m_i \in$

$o\text{-ints}_{\mathcal{I}}(u_i)$, $1 \leq i \leq k$, be an o -interpretation at u_i . Then the join $m_1 \bowtie \cdots \bowtie m_k = \bigcup_{1 \leq i \leq k} \text{int}(m_i)$ at u is defined iff for each $u' \in \text{fai}(u)$ the set of o -interpretations at u' that are reachable (in \bar{F}) from some o -interpretation m_i , $1 \leq i \leq k$, contains exactly one o -interpretation $m' \in o\text{-ints}_{\mathcal{I}}(u')$.

Intuitively, a set of interpretations can only be joined if all interpretations depend on the same (and on a single) interpretation at every unit.

Example 25 (ctd) In \mathcal{I}_2 , i -interpretations m_1, m_4, m_5, m_7 , and m_8 are created by trivial join operations with none or one predecessor unit. For m_{13} and m_{14} , we have a nontrivial join: $\text{int}(m_{13}) = \text{int}(m_6) \cup \text{int}(m_{11})$ and the join is defined because $\text{fai}(u_4) = \{u_1\}$, and from m_6 and m_{11} we can reach in \mathcal{I}_2 exactly one o -interpretation at u_1 . Observe that the join $m_6 \bowtie m_9$ is not defined, as we can reach in \mathcal{I}_2 from $\{m_6, m_9\}$ the o -interpretations m_2 and m_3 at u_1 , and thus more than exactly one o -interpretation at some FAI of u_4 . Similarly, the join $m_6 \bowtie m_{10}$ is undefined, as we can reach m_2 and m_3 at u_1 . \square

The result of a join is the union of predecessor interpretations; this is important for answer set graphs and join operations on them, which comes next. Note that each leaf unit (i.e., without predecessors) has exactly one well-defined join result, viz. \emptyset .

If we add a new i -interpretation from the result of a join operation to an i -graph and dependencies to all participating o -interpretations, the resulting graph is again an i -graph; thus the join is sound wrt. to the i -graph properties. Moreover, each i -interpretation that can be added to a i -graph while preserving the i -graph conditions can be synthesized by a join; that is, the join is complete for such additions. This is a consequence of the following result.

Proposition 4 Let $\mathcal{I} = (M, F, \text{unit}, \text{type}, \text{int})$ be an i -graph for an evaluation graph $\mathcal{E} = (V, E)$ and $u \in V$ with $\text{preds}_{\mathcal{E}}(u) = \{u_1, \dots, u_k\}$. Furthermore, let $m_i \in o\text{-ints}_{\mathcal{I}}(u_i)$, $1 \leq i \leq k$, such that no vertex $m \in i\text{-ints}_{\mathcal{I}}(u)$ exists such that $\{(m, m_1), \dots, (m, m_k)\} \subseteq F$. Then the join $J = m_1 \bowtie \cdots \bowtie m_k$ is defined at u iff $\mathcal{I}' = (M', F', \text{unit}', \text{type}', \text{int}')$ is an i -graph for \mathcal{E} where (a) $M' = M \cup \{m'\}$ for some new vertex $m' \in \mathcal{I}_{id} \setminus M$, (b) $F' = F \cup \{(m', m_i) \mid 1 \leq i \leq k\}$, (c) $\text{unit}' = \text{unit} \cup \{(m', u)\}$, (d) $\text{type}' = \text{type} \cup \{(m', \text{I})\}$, and (e) $\text{int}' = \text{int} \cup \{(m', J)\}$.

Note that the i -graph definition specifies topological properties of an i -graph wrt. an evaluation graph. In the following we extend this specification to the contents of interpretations.

5.3 Answer Set Graph

We next restrict i -graphs to *answer set graphs* such that interpretations correspond with answer sets of certain HEX programs that are induced by the evaluation graph.

Definition 19 (Answer Set Graph) An answer set graph $\mathcal{A} = (M, F, \text{unit}, \text{type}, \text{int})$ for an evaluation graph $\mathcal{E} = (U, E)$ is an i -graph for \mathcal{E} such that for each unit $u \in U$, it holds that

- (a) $\{\text{int}^+(m) \mid m \in i\text{-ints}_{\mathcal{I}}(u)\} \subseteq \mathcal{AS}(u^<)$, i.e., every expanded i -interpretation at u is an answer set of $u^<$;
- (b) $\{\text{int}^+(m) \mid m \in o\text{-ints}_{\mathcal{I}}(u)\} \subseteq \mathcal{AS}(u^{\leq})$, i.e., every expanded o -interpretation at u is an answer set of u^{\leq} ; and
- (c) for each $m \in i\text{-ints}_{\mathcal{I}}(u)$, it holds that $\text{int}(m) = \bigcup_{(m, m_i) \in F} \text{int}(m_i)$.

Note that each leaf unit u , has $u^< = \emptyset$, and thus \emptyset is the only i-interpretation possible. Moreover, condition (c) is necessary to ensure that an i-interpretation at unit u contains all atoms of answer sets of predecessor units that are relevant for evaluating u . Furthermore, note that the empty graph is an answer set graph.

Example 26 (ctd) *The example i-graph \mathcal{I}_2 is in fact an answer set graph. First, $int^+(m_1) = \emptyset$ and $u_1^< = \emptyset$ and indeed $\emptyset \in \mathcal{AS}(\emptyset)$ which satisfies condition (a). Less obvious is the case of o-interpretation m_6 in \mathcal{I}_2 : $int^+(m_6) = \{swim(out)\}$ and $u_2^< = \{r_1, r_2, c_8\}$; as c_8 kills all answer sets where money is required, $\mathcal{AS}(\{r_1, r_2, c_8\}) = \{\{swim(out)\}\}$; hence $int^+(m_6)$ is the only expanded interpretation of an o-interpretation possible at u_2 . Furthermore, the condition (IG-U) on i-graphs implies that m_6 is the only possible o-interpretation at u_2 . Consider next m_{13} :*

$$u_4^< = \{r_1, r_2, r_3, r_4, c_6, c_7, c_8\} \text{ and} \\ int^+(m_{13}) = \{go, goto(altD), ngoto(gansD), swim(out)\}.$$

The two answer sets of $u_4^<$ are $\{go, goto(altD), ngoto(gansD), swim(out)\}$, and $\{go, goto(gansD), ngoto(altD), swim(out)\}$, and $int^+(m_{13})$ is one of them; the other one is $int^+(m_{14})$. Finally

$$int^+(m_{15}) = \{swim(out), goto(altD), go, ngoto(gansD), need(loc, yogamat)\},$$

which is the single answer set of $u_4^< = P_{swim}$. \square

Similarly as for i-graphs, the join is a sound and complete operation to add i-interpretations to an answer set graph.

Proposition 5 *Let $\mathcal{A} = (M, F, unit, type, int)$ be an answer set graph for an evaluation graph $\mathcal{E} = (V, E)$ and let $u \in V$ with $preds_{\mathcal{E}}(u) = \{u_1, \dots, u_k\}$. Furthermore, let $m_i \in o-ints_{\mathcal{A}}(u_i)$, $1 \leq i \leq k$, such that no $m \in i-ints_{\mathcal{A}}(u)$ with $\{(m, m_1), \dots, (m, m_k)\} \subseteq F$ exists. Then the join $J = m_1 \bowtie \dots \bowtie m_k$ is defined at u iff $\mathcal{A}' = (M', F', unit', type', int')$ is an answer set graph for \mathcal{E} where (a) $M' = M \cup \{m'\}$ for some new vertex $m' \in \mathcal{I}_{id} \setminus M$, (b) $F' = F \cup \{(m', m_i) \mid 1 \leq i \leq k\}$, (c) $unit' = unit \cup \{(m', u)\}$, (d) $type' = type \cup \{(m', 1)\}$, and (e) $int' = int \cup \{(m', J)\}$.*

Example 27 (ctd) *Imagine that \mathcal{I}_2 has no interpretations at u_4 . The following candidate pairs of o-interpretations exist for creating i-interpretations at u_4 : $m_6 \bowtie m_9$, $m_6 \bowtie m_{10}$, $m_6 \bowtie m_{11}$, and $m_6 \bowtie m_{12}$. As seen in Example 25, $m_{13} = m_6 \bowtie m_{11}$ and $m_{14} = m_6 \bowtie m_{12}$ are the only joins at u_4 that are defined. In Example 26 we have seen that $\mathcal{AS}(u_4^<) = \{int^+(m_{13}), int^+(m_{14})\}$, and due to (IG-U), we cannot have additional i-interpretations with the same content. \square*

5.3.1 Complete Answer Set Graphs

We next introduce a notion of completeness for answer set graphs.

Definition 20 *Let $\mathcal{A} = (M, F, unit, type, int)$ be an answer set graph for an evaluation graph $\mathcal{E} = (U, E)$ and let $u \in U$. Then*

- \mathcal{A} is input-complete for u , if $\{int^+(m) \mid m \in i-ints_{\mathcal{A}}(u)\} = \mathcal{AS}(u^<)$, and
- \mathcal{A} is output-complete for u , if $\{int^+(m) \mid m \in o-ints_{\mathcal{A}}(u)\} = \mathcal{AS}(u^{\leq})$.

If an answer set graph is complete for all units of its corresponding evaluation graph, answer sets of the associated program can be obtained as follows.

Theorem 5 *Let $\mathcal{E} = (U, E)$, where $U = \{u_1, \dots, u_n\}$, be an evaluation graph of a program P , and let $\mathcal{A} = (M, F, \text{unit}, \text{type}, \text{int})$ be an answer set graph that is output-complete for every unit $u \in U$. Then*

$$\mathcal{AS}(P) = \left\{ \bigcup_{i=1}^n \text{int}(m_i) \mid m_i \in o\text{-ints}_{\mathcal{A}}(u_i), 1 \leq i \leq n, |o\text{-ints}_{\mathcal{A}'}(u_i)| = 1 \right\}, \quad (4)$$

where \mathcal{A}' is the subgraph of \mathcal{A} consisting of all interpretations that are reachable in \mathcal{A} from some interpretation m_1, \dots, m_n .

Example 28 (ctd) *In \mathcal{I}_2 we first choose $m_{15} \in o\text{-ints}(u_4)$, which is the only o-interpretation at u_4 . The subgraph reachable from m_{15} must contain exactly one o-interpretation at each unit; we thus must choose every o-interpretations m such that $m_{15} \rightarrow^+ m$. Hence we obtain*

$$\begin{aligned} & \{ \text{int}(m_3) \cup \text{int}(m_6) \cup \text{int}(m_{11}) \cup \text{int}(m_{15}) \} \\ &= \{ \{ \text{swim}(\text{out}) \} \cup \emptyset \cup \{ \text{goto}(\text{altD}), \text{ngoto}(\text{gansD}), \text{go} \} \cup \{ \text{need}(\text{loc}, \text{yogamat}) \} \} \\ &= \{ \{ \text{swim}(\text{out}), \text{goto}(\text{altD}), \text{ngoto}(\text{gansD}), \text{go}, \text{need}(\text{loc}, \text{yogamat}) \} \} \end{aligned}$$

which is indeed the set of answer sets of P_{swim} . \square

The rather involved set construction in (4) establishes a relationship between answer sets of a program and complete answer set graphs that resembles condition (IG-F) of i-graphs. To obtain a more convenient way to enumerate answer sets, we can extend an evaluation graph always with a single void unit u_{final} that depends on all other units in the graph (i.e., $(u_{\text{final}}, u) \in E$ for each $u \in U \setminus \{u_{\text{final}}\}$), which we call a *final unit*; the answer sets of P correspond then directly to i-interpretations at u_{final} . Formally,

Proposition 6 *Let $\mathcal{A} = (M, F, \text{unit}, \text{type}, \text{int})$ be an answer set graph for an evaluation graph $\mathcal{E} = (U, E)$ of a program P , where \mathcal{E} contains a final unit u_{final} , and assume that \mathcal{A} is input-complete for U and output-complete for $U \setminus \{u_{\text{final}}\}$. Then*

$$\mathcal{AS}(P) = \{ \text{int}(m) \mid m \in i\text{-ints}_{\mathcal{A}}(u_{\text{final}}) \}. \quad (5)$$

Expanding i-interpretations at u_{final} is not necessary, as u_{final} depends on all other units; thus for every $m \in i\text{-ints}_{\mathcal{A}}(u_{\text{final}})$ it holds that $\text{int}^+(m) = \text{int}(m)$.

We will use the technique with u_{final} for our model enumeration algorithm; as the join condition must be checked anyways, this technique is an efficient and simple method for obtaining all answer sets of a program using an answer set graph.

5.4 Answer Set Building

Thanks to the results above, we can obtain the answer sets of a HEX-program from any answer set graph for it. To build an answer set graph, we proceed as follows. We start with an empty graph, obtain o-interpretations by evaluating a unit on an i-interpretation, and then gradually generate i-interpretations by joining o-interpretations of predecessor units in an evaluation graph at hand.

Towards an algorithm for evaluating a HEX-program based on an evaluation graph, we use a generic grounding algorithm GROUNDHEX for lde-safe programs, and a solving algorithm EVALUATEGROUNDHEX which returns for a ground HEX-program P its answer sets $\mathcal{AS}(P)$. We assume that they satisfy the following properties.

Algorithm 1: EVALUATELDES SAFE

Input: A liberally de-safe HEX-program P , an input interpretation I
Output: All answer sets of $P \cup \text{facts}(I)$ without I
 // add input facts and ground, cf. (15)
 $P' \leftarrow \text{GROUNDHEX}(P \cup \text{facts}(I))$
 // evaluate the ground program, cf. (16),
 // and perform output projection
 return $\{I' \setminus I \mid I' \in \text{EVALUATEGROUNDHEX}(P')\}$

Algorithm 2: BUILDANSWERSETS

Input: $\mathcal{E} = (V, E)$: evaluation graph for HEX program P , which contains a unit u_{final} that depends on all other units in V
Output: a set of all answer sets of P
 $M := \emptyset, F := \emptyset, \text{unit} := \emptyset, \text{type} := \emptyset, \text{int} := \emptyset, U := V$

```

(a) while  $U \neq \emptyset$  do
    choose  $u \in U$  s.t.  $\text{preds}_{\mathcal{E}}(u) \cap U = \emptyset$ 
    let  $\{u_1, \dots, u_k\} = \text{preds}_{\mathcal{E}}(u)$ 
    if  $k = 0$  then
      (b)  $m := \max(M) + 1$ 
           $M := M \cup \{m\}$ 
           $\text{unit}(m) := u, \text{type}(m) := \text{I}, \text{int}(m) := \emptyset$ 
    else
      (c) for  $m_1 \in o\text{-ints}(u_1), \dots, m_k \in o\text{-ints}(u_k)$  do
          if  $J = m_1 \bowtie \dots \bowtie m_k$  is defined then
             $m := \max(M) + 1$ 
             $M := M \cup \{m\}, F := F \cup \{(m, m_i) \mid 1 \leq i \leq k\}$ 
             $\text{unit}(m) := u, \text{type}(m) := \text{I}, \text{int}(m) := J$ 
      (d) if  $u = u_{\text{final}}$  then
          return  $i\text{-ints}(u_{\text{final}})$ 
      (e) for  $m' \in i\text{-ints}(u)$  do
           $O := \text{EVALUATELDES SAFE}(u, \text{int}(m'))$ 
          for  $o \in O$  do
             $m := \max(M) + 1$ 
             $M := M \cup \{m\}, F := F \cup \{(m, m')\}$ 
             $\text{unit}(m) := u, \text{type}(m) := O, \text{int}(m) := o$ 
      (f)  $U := U \setminus \{u\}$ 

```

Property 1 Given an lde-safe program P , $\text{GROUNDHEX}(P)$ returns a finite ground program such that $\mathcal{AS}(P) = \mathcal{AS}(\text{GROUNDHEX}(P))$.

Property 2 Given a finite ground HEX-program P , $\text{EVALUATEGROUNDHEX}(P) = \mathcal{AS}(P)$.

Concrete such algorithms are given in (15) and (16), respectively. The idea of the grounding algorithm is to evaluate the external atoms in the program under a (finite) number of relevant inputs in order to determine the relevant set of constants in advance, while the solving algorithm is based on conflict-driven clause learning (CDCL) and lifts the work of (25) from ordinary to HEX programs.

By composing the two algorithms, we obtain Algorithm 1 for evaluating a single unit. Formally, it has the following property.

Proposition 7 Given an lde-safe HEX-program P and an input interpretation I , Algorithm 1 returns the set $\{I' \setminus I \mid I' \in \mathcal{AS}(P \cup \text{facts}(I))\}$, i.e., the answer sets of P augmented with facts for the input I , projected to the non-input.

We are now ready to formulate an algorithm for evaluating HEX programs that have been decomposed into an evaluation graph.

To this end, we build first an evaluation graph \mathcal{E} and then compute gradually an answer set graph $\mathcal{A} = (M, F, \text{unit}, \text{type}, \text{int})$ based on \mathcal{E} , proceeding along already evaluated units towards the unit u_{final} . Algorithm 2 shows the model building algorithm in pseudo-code, in which the positive integers $\mathbb{N} = \{1, 2, \dots\}$

are used as identifiers \mathcal{I}_{id} and $\max(M)$ is maximum in any set $M \subseteq \mathbb{N}$ where, by convention, $\max(\emptyset) = 0$. Intuitively, the algorithm works as follows. The set U contains units for which \mathcal{A} is not yet output-complete (see Definition 20); we start with an empty answer set graph \mathcal{A} , thus initially $U = V$. In each iteration of the while loop (a), a unit u that is not output-complete and depends only on output-complete units is selected. The first for loop (c) makes u input-complete; if u is the final unit, the answer sets are returned in (d), otherwise the second for loop (e) makes u output-complete, and then u is removed from U . Each iteration makes one unit input- and output-complete; hence when the algorithm reaches u_{final} and makes it input-complete, all answer sets can directly be returned in (d). Formally, we have

Theorem 6 *Given an evaluation graph $\mathcal{E} = (V, E)$ of a HEX program P , BUILDANSWERSETS(\mathcal{E}) returns $\mathcal{AS}(P)$.*

A run of the algorithm on our running example using the evaluation graph \mathcal{E}_2 extended with a final unit is given in Appendix B.

5.4.1 Model Streaming

Algorithm BUILDANSWERSETS as described above keeps all answer sets in memory, and it evaluates each unit only once wrt. every possible i -interpretation. This may lead to a resource bound excess, as in general an exponential number of answer sets respectively interpretations at evaluation units are possible. However, keeping the whole answer set graph in memory is not necessary for computing all answer sets.

We have realized a variant of Algorithm BUILDANSWERSETS that uses the same principle of constructing an answer set graph, interpretations are created at a unit *on demand* when they are requested by units that depend on it; furthermore, the algorithm keeps basically only one interpretation at each evaluation unit in memory at a time, which means that interpretations are provided in a *streaming fashion* one by one, and likewise the answer sets of the program at the unit u_{final} , where the model building starts. Such answer set streaming is particularly attractive for applications, as one can terminate the computation after obtaining sufficiently many answer sets. On the other hand, it comes at the cost of potential re-evaluation of units wrt. the same i -interpretation, as we need to trade space for time. However, in practice this algorithm works well and is the one used in the dlhex prototype. We describe this algorithm in Appendix C.

6 Implementation

In this section we give some details on the implementation of the techniques. Our prototype system is called dlhex; it is written in C++ and online available as open-source software.⁷ The current version 2.4.0 was released in September 2014.

We first describe the general architecture, the major components, and their interplay. Then we give an overview about the existing heuristics. For details on the usage of the system, we refer to the website; an exhaustive description of the supported command-line parameters is output when the system is called without parameters.

6.1 System Architecture

The dlhex system architecture is shown in Figure 9. The arcs model both control and data flow within the system. The evaluation of a HEX-program works as follows.

⁷<http://www.kr.tuwien.ac.at/research/systems/dlhex>

First, the input program is read from the file system or from standard input and passed to the *evaluation framework* ①. The evaluation framework creates then an *evaluation graph* depending on the chosen evaluation heuristics. This results in a number of interconnected *evaluation units*. While the interplay of the units is managed by the evaluation framework, the individual units are handled by *model generators* of different kinds.

Each instance of a model generator takes care of a single evaluation unit, receives *input interpretations* from the framework (which are either output by predecessor units or come from the input facts for leaf units), and sends output interpretations back to the framework ②, which manages the integration of the latter to final answer sets.

Internally, the model generators make use of a *grounder* and a *solver* for ordinary ASP programs. The architecture of our system is flexible and supports multiple concrete backends that can be plugged in. Currently it supports *dlv*, *gringo* 4.4.0 and *clasp* 3.1.0, as well as an internal grounder and a solver that were built from scratch (mainly for testing purposes); they use basically the same core algorithms as *gringo* and *clasp*, but without optimizations. The reasoner backends *gringo* and *clasp* are statically linked to our system; thus no interprocess communication is necessary. The model generator within the *dlvhex* core sends a non-ground program to the HEX-grounder, and receives a ground program ③. The HEX-grounder in turn uses an ordinary ASP grounder as submodule ④ and accesses external sources to handle value invention ⑤. The ground-program is then sent to the solver and answer sets of the ground program (i.e. candidate compatible sets) are returned ⑥. Note that the grounder and the solver are separated and communicate only via the model generator; this is in contrast to previous *dlvhex* versions, where the external grounder and solver formed a single unit (i.e., the non-ground program was sent and the answer sets were retrieved). Separating the two units became necessary as the *dlvhex* core needs access to the ground program; otherwise important structural information, e.g. cyclicity, would be hidden.

The solver backend makes callbacks to the *post propagator* in the *dlvhex* core once a model has been found or after unit and unfounded set propagation has been finished (actually, with *dlv* backend callbacks occur only after a candidate compatible set has been found, but not during model building). During the callback, a complete or partial model is sent from the solver backend to the post propagator, and learned nogoods are sent back to the external solver ⑦. For *clasp* as backend, we exploit its SMT interface, which was previously used for the special case of constraint answer set solving (22). The post propagator has then two key tasks: *compatibility checking with learning* and *unfounded set detection*. Compatibility checking, as formalized by 14), checks whether the guesses of the external atom replacements by the ordinary ASP solver coincide with the actual values of the external source. This requires calls to the *plugins* implementing the external sources. The input list is sent to the external source, and the truth values and possibly learned nogoods are returned to the post propagator ⑧. Moreover, the post propagator also sends the (complete or partial) model to the *unfounded set checker (UFS checker)* to find unfounded sets that are not detected by the ordinary ASP solver (i.e., caused by external sources). The UFS checker employs a SAT solver ⑩, which can either be *clasp* or the internal solver, and possibly returns nogoods learned from unfounded sets to the post propagator ⑧. UFS detection also needs to call the external sources for guess verification, as shown by 16) ⑩. The post propagator sends all learned nogoods back to the ASP solver. This ensures that eventually only valid answer sets arrive at the model generator ⑥.

Finally, after the evaluation framework has built the final answer sets from the output interpretations of the individual evaluation units, they are output to the user ⑫.

Example 29 *The program from Example 2 is encoded in file `swimming.hex` as follows, where as in the *dlv* language `:-` stands for \leftarrow and the letter \vee for \vee .*

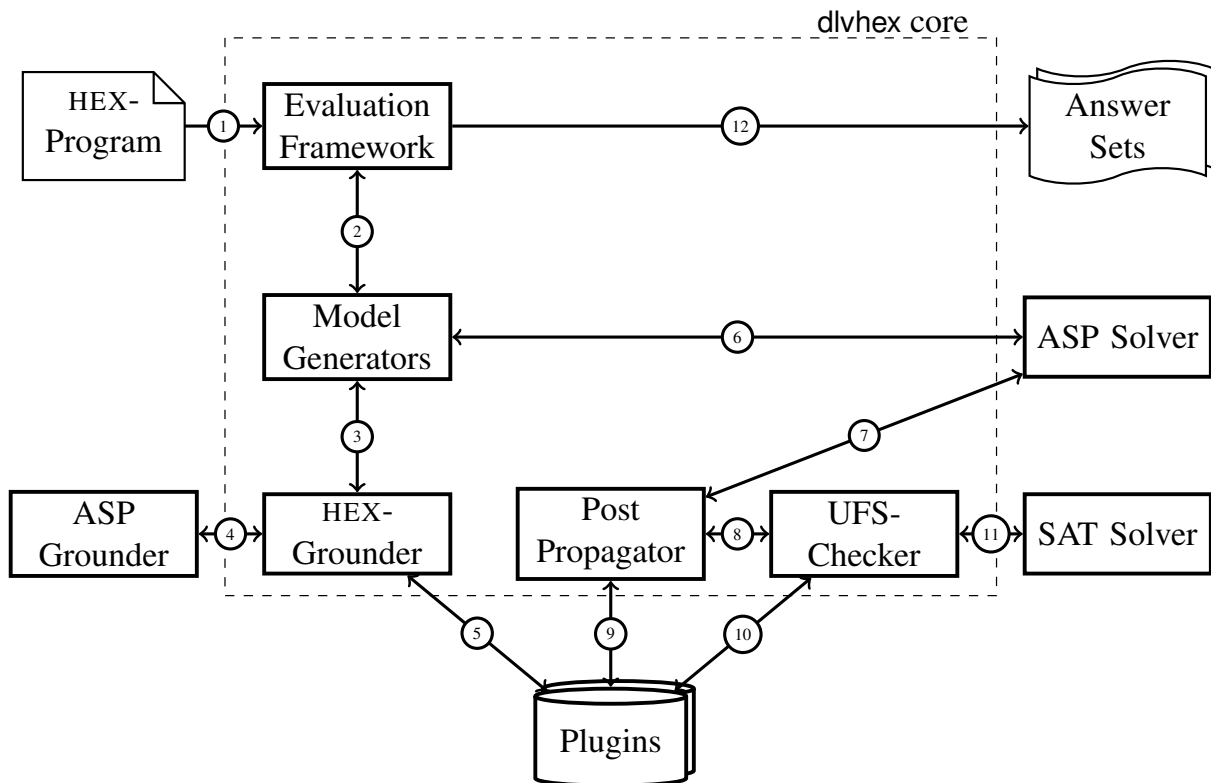


Figure 9: Architecture of dlhex

```

swim(in) v swim(out).
need(inout, C) :- &rq[swim](C).
goto(X) v ngoto(C) :- swim(P), location(P, C).
go :- goto(X).
need(loc, C) :- &rq[goto](C).
:- goto(X), goto(Y), X != Y.
:- not go.
:- need(X, money).

```

6.2 Heuristics

As for creating evaluation graphs, several heuristics have been implemented. A heuristics starts with the rule dependency graph as by Definition 10 and then acyclically combines nodes into units.

Some heuristics are described in the following.

H0 is a ‘trivial’ heuristics that makes units as small as possible, which is useful for debugging, however this generates the largest possible number of evaluation units and therefore incurs a large overhead.

H1 is the evaluation heuristics of the dlhex prototype version 1. **H1** makes units as large as possible and has several drawbacks as discussed above.

H2 is a simple evaluation heuristics which has the goal of finding a compromise between the **H0** and **H1**.

It places rules into units as follows:

- (i) it puts rules r_1, r_2 into the same unit whenever $r_1 \rightarrow_{m,n} s$ and $r_2 \rightarrow_{m,n} s$ for some rule s and there is no rule t such that exactly one of r_1, r_2 depends on t ;
- (ii) it puts rules r_1, r_2 into the same unit whenever $s \rightarrow_{m,n} r_1$ and $s \rightarrow_{m,n} r_2$ for some rule s and there is no rule t such that t depends on exactly one of r_1, r_2 ; but
- (iii) it never puts rules r, s into the same unit if r contains external atoms and $r \rightarrow_{m,n} s$.

Intuitively, *H2* builds an evaluation graph that puts all rules with external atoms and their successors into one unit, while separating rules creating input for distinct external atoms. This avoids redundant computation and joining unrelated interpretations.

H3 is a heuristics for generating a good generalized evaluation graph by iteratively merging units. It aims at two opposing goals: (1) minimizing the number of units, and (2) splitting the program whenever a de-relevant nonmonotonic external atom would receive input from the same unit. *H3* greedily gives preference to (1) and is motivated by the following considerations.

The grounding algorithm by 15) evaluates the external sources under all interpretations such that the set of observed constants is maximized. While monotonic and antimonotonic input atoms are not problematic (the algorithm can simply set all to true resp. false), nonmonotonic parameters require an exponential number of evaluations in general. Thus, although program decomposition is not strictly necessary, it is still useful in such cases as it restricts grounding to those interpretations that are actually relevant in some answer set. However, on the other hand it can be disadvantageous for propositional solving algorithms such as those in (14). Program decomposition can be seen as a hybrid between traditional and lazy grounding (cf. e.g. 41)), as program parts are instantiated that are larger than single rules but smaller than the whole program.

6.3 Experimental Results

In this section, we evaluate the model-building framework empirically. To this end, we compare the following configurations. In the *old* column, we use the previous evaluation method before the framework was developed (48). This method also makes use of program decomposition. However, in contrast to our new framework the decomposition is based on atom dependencies rather than rule dependencies, and the decomposition strategy is hard-coded and not customizable. This evaluation method corresponds to heuristics *H1* in our new framework and can thus be emulated.

In the *w/o framework* column, we present the results without application of the framework using the grounding algorithm by 15). Note that before this grounding algorithm was developed, a direct evaluation was not possible since program decomposition was necessary for grounding purposes (at that time using the strategy by 48) based on atom dependencies). With the new grounding algorithm, decomposition is not necessary anymore, but still useful as the results in the *new* column show, where we present the results when the default heuristics of the new framework is used.

The configuration of the grounding algorithm and the solving algorithm (e.g. conflict-driven learning strategies) also influence the results. Moreover, in addition to the default heuristics of framework, other heuristics have been developed as well and the best selection of the heuristics often depends on the configuration of the grounding and the solving algorithm. Since they were used as black boxes in Algorithm 1, an exhaustive experimental analysis of the system is beyond the scope of this paper and would require an in-depth description of these algorithms. Thus, we confine the discussion to the default settings, which suffices to show that the new framework can speed up the evaluation significantly. For an in depth discussion,

Topology and Instance Size	First Answer Set			All Answer Sets		
	old	w/o framework	new	old	w/o framework	new
d-7-7-3-3 (10)	1.23 (0)	0.29 (0)	0.38 (0)	4.93 (0)	0.76 (0)	0.79 (0)
d-7-7-4-4 (10)	18.43 (0)	1.09 (0)	0.76 (0)	50.78 (0)	3.39 (0)	1.80 (0)
d-7-7-5-5 (10)	94.18 (1)	3.60 (0)	1.52 (0)	289.35 (4)	20.21 (0)	4.97 (0)
h-9-9-3-3 (10)	83.17 (1)	3.77 (0)	0.70 (0)	300.96 (4)	28.67 (0)	2.11 (0)
h-9-9-4-4 (10)	389.74 (6)	30.56 (0)	2.14 (0)	555.94 (9)	335.11 (5)	12.56 (0)
r-7-7-4-4 (10)	39.27 (0)	2.82 (0)	0.33 (0)	366.17 (5)	57.26 (0)	2.06 (0)
r-7-7-5-5 (10)	389.88 (6)	105.80 (1)	0.93 (0)	600.00 (10)	377.37 (5)	4.39 (0)
r-7-8-5-5 (10)	226.04 (3)	25.11 (0)	0.57 (0)	541.80 (9)	317.64 (5)	3.99 (0)
r-7-9-5-5 (10)	355.37 (5)	145.99 (2)	0.87 (0)	600.00 (10)	458.14 (7)	5.42 (0)
r-8-7-5-5 (10)	502.64 (8)	329.47 (5)	1.21 (0)	555.26 (9)	443.15 (7)	5.84 (0)
r-8-8-5-5 (10)	390.81 (6)	201.08 (3)	1.00 (0)	600.00 (10)	495.41 (8)	5.38 (0)
z-7-7-3-3 (10)	2.34 (0)	0.32 (0)	0.44 (0)	9.17 (0)	1.13 (0)	1.00 (0)
z-7-7-4-4 (10)	33.32 (0)	1.58 (0)	1.07 (0)	182.44 (2)	9.00 (0)	2.67 (0)
z-7-7-5-5 (10)	164.33 (2)	12.69 (0)	3.52 (0)	502.49 (8)	89.01 (1)	6.90 (0)

Table 1: MCS experiments: variable topology (d, h, r, z) and instance size.

we refer to 16 (16; 15) and 46), where the efficiency was evaluated using a variety of applications including planning tasks (e.g. robots searching an unknown area for an object, tour planning), computing extensions of abstract argumentation frameworks, inconsistency analysis in multi-context systems, and reasoning over description logic knowledge bases.

We discuss here two benchmark problems, which we evaluated on a Linux server with two 12-core AMD 6176 SE CPUs with 128GB RAM running an *HTCondor* load distribution system⁸ that ensures robust runtimes (i.e., multiple runs of the same instance have negligible deviations) and *dlvhex* version 2.4.0. The grounder and solver backends for all benchmarks are *gringo* 4.4.0 and *clasp* 3.1.1. For each instance, we limited the CPU usage to two cores and 8GB RAM. The timeout for each instance was 600 seconds. Each line shows the average runtimes over all instances of a certain size, where each timeout counts as 600 seconds. Numbers in parentheses are the numbers of instances of respective size in the leftmost column and the numbers of timeout instances elsewhere. The generators, instances and external sources are available at <http://www.kr.tuwien.ac.at/research/projects/hexhex/hexframework>.

6.3.1 Multi-Context Systems (MCS)

The MCS benchmarks originate in the application scenario of enumerating output-projected equilibria (i.e., global models) of a given multi-context system (MCS) (cf. Section 2.3.2). Each instance comprises 7–9 contexts (propositional knowledge bases) whose local semantics is modeled by external atoms; roughly speaking, they single out assignments to the atoms of a context occurring in bridge rules such that local models exist. For each context, 5–10 such atoms are guessed and bridge rules, which are modeled by ordinary rules, are randomly constructed on top. The MCS instances were generated using the DMCS (2) instance generator, with 10 randomized instances for different link structure between contexts (diamond (d), house (h), ring (r), zig-zag (z)) and system size; they have between 4 and about 20,000 answer sets, with an average of 400. We refer to (2) and (49) for more details on the benchmarks and the HEX-programs.

Table 1 shows the experimental results: computation with the old method often exceeds the time limit, while the new method *H2* manages to enumerate all solutions of all instances. Monolithic evaluation without decomposition shows a performance between the old and new method. These results show that our new evaluation method is essential for using HEX to computationally realize the MCS application.

⁸<http://research.cs.wisc.edu/htcondor>

Instance Size	First Answer Set			All Answer Sets		
	old	w/o framework	new	old	w/o framework	new
1 (1)	2.84 (0)	3.14 (0)	2.78 (0)	2.73 (0)	3.14 (0)	2.79 (0)
2 (1)	6.13 (0)	7.18 (0)	4.90 (0)	6.05 (0)	7.17 (0)	4.88 (0)
3 (1)	10.18 (0)	12.30 (0)	8.32 (0)	10.25 (0)	12.35 (0)	8.37 (0)
4 (1)	15.92 (0)	18.66 (0)	12.12 (0)	15.86 (0)	18.85 (0)	12.16 (0)
5 (1)	26.06 (0)	28.47 (0)	17.17 (0)	26.23 (0)	28.35 (0)	17.06 (0)
6 (1)	47.06 (0)	45.71 (0)	23.39 (0)	46.84 (0)	45.62 (0)	23.26 (0)
7 (1)	92.76 (0)	79.41 (0)	31.19 (0)	96.56 (0)	79.82 (0)	31.04 (0)
8 (1)	198.59 (0)	155.10 (0)	37.85 (0)	199.74 (0)	155.26 (0)	38.06 (0)
9 (1)	600.00 (1)	600.00 (1)	46.61 (0)	600.00 (1)	600.00 (1)	46.75 (0)
10 (1)	600.00 (1)	600.00 (1)	57.48 (0)	600.00 (1)	600.00 (1)	57.40 (0)
11 (1)	600.00 (1)	600.00 (1)	68.98 (0)	600.00 (1)	600.00 (1)	69.45 (0)
12 (1)	600.00 (1)	600.00 (1)	84.41 (0)	600.00 (1)	600.00 (1)	84.11 (0)
13 (1)	600.00 (1)	600.00 (1)	99.55 (0)	600.00 (1)	600.00 (1)	99.52 (0)
14 (1)	600.00 (1)	600.00 (1)	117.39 (0)	600.00 (1)	600.00 (1)	117.15 (0)
15 (1)	600.00 (1)	600.00 (1)	138.45 (0)	600.00 (1)	600.00 (1)	137.51 (0)
16 (1)	600.00 (1)	600.00 (1)	163.12 (0)	600.00 (1)	600.00 (1)	158.43 (0)
17 (1)	600.00 (1)	600.00 (1)	184.99 (0)	600.00 (1)	600.00 (1)	181.94 (0)
18 (1)	600.00 (1)	600.00 (1)	208.83 (0)	600.00 (1)	600.00 (1)	210.82 (0)
19 (1)	600.00 (1)	600.00 (1)	236.98 (0)	600.00 (1)	600.00 (1)	237.45 (0)
20 (1)	600.00 (1)	600.00 (1)	267.54 (0)	600.00 (1)	600.00 (1)	268.60 (0)
21 (1)	600.00 (1)	600.00 (1)	600.00 (1)	600.00 (1)	600.00 (1)	600.00 (1)

Table 2: RSTRACK experiments: variable number of conference tracks, single answer set.

6.3.2 Reviewer Selection (RS)

Our second benchmark is Reviewer Selection (RS): we represent c conference tracks, r reviewers and p papers. Papers and reviewers are assigned to conference tracks, and there are conflicts between reviewers and papers, some of which are given by external atoms. We consider two scenarios: RSTRACK and RSPAPER. They are designed to measure the effect of external atoms on the elimination of a large number of answer set candidates; in contrast to the MCS experiments we can control this aspect in the RS experiments.

In RSTRACK we vary the number c of conference tracks, where each track has 20 papers and 20 reviewers. Each paper must get two reviews, and no reviewer must get more than two papers. Conflicts are dense such that only one valid assignment exists per track, hence each instance has exactly one answer set, and in each track two conflicts are external. For each number c there is only one instance because RSTRACK instances are not randomized. The results of RSTRACK are shown in Table 2: runtimes of the old evaluation heuristics ($H1$) grow fastest with size, without using decomposition grows slightly slower but also reaches timeout at size 9. Only the new decomposition ($H2$ heuristics) can deal with size 20 without timeout. Finding the first answer set and enumerating all answer sets show very similar times, as RSTRACK instances have a single answer set and finding it seems hard.

In RSPAPER we fix the number of tracks to $c=5$; we vary the number p of papers in each track and set the number of reviewers to $r=p$. Each paper must get three reviews and each reviewer must not get more than three papers assigned. Conflicts are randomized and less dense than in RSTRACK: the number of answer sets is greater than one and does not grow with the instance size. Over all tracks and papers, $2p$ randomly chosen conflicts are external, and we generate 10 random instances per size and report results averaged per instance size in Table 3. As clearly seen, our new method is always faster than the other methods, and evaluation without a decomposition framework performs slightly better than the old method. Different from RSTRACK, we can see a clear difference between finding the first answer set and enumerating all answer sets as RSPAPER instances have more than one answer set.

To confirm that the new method is geared towards handling many external atoms, we conducted also

Instance Size	First Answer Set			All Answer Sets		
	old	w/o framework	new	old	w/o framework	new
5 (10)	1.06 (0)	0.28 (0)	0.21 (0)	2.25 (0)	0.43 (0)	0.23 (0)
8 (10)	8.76 (0)	2.73 (0)	0.38 (0)	14.73 (0)	4.54 (0)	0.44 (0)
11 (10)	108.70 (1)	83.26 (1)	0.98 (0)	171.01 (2)	104.84 (1)	1.28 (0)
14 (10)	180.99 (2)	125.83 (1)	2.08 (0)	299.22 (4)	245.62 (3)	2.67 (0)
17 (10)	418.92 (6)	364.95 (5)	5.15 (0)	549.01 (9)	513.21 (8)	8.14 (0)
20 (10)	485.35 (8)	453.39 (7)	7.32 (0)	507.66 (8)	501.74 (8)	14.45 (0)
23 (10)	542.03 (9)	508.75 (8)	13.91 (0)	600.00 (10)	600.00 (10)	23.16 (0)
26 (10)	600.00 (10)	600.00 (10)	33.20 (0)	600.00 (10)	600.00 (10)	154.51 (2)
29 (10)	600.00 (10)	600.00 (10)	60.78 (0)	600.00 (10)	600.00 (10)	108.03 (0)
32 (10)	600.00 (10)	600.00 (10)	129.95 (0)	600.00 (10)	600.00 (10)	315.56 (4)
35 (10)	600.00 (10)	600.00 (10)	136.84 (0)	600.00 (10)	600.00 (10)	302.90 (3)
38 (10)	600.00 (10)	600.00 (10)	308.92 (3)	600.00 (10)	600.00 (10)	441.06 (6)
41 (10)	600.00 (10)	600.00 (10)	421.69 (6)	600.00 (10)	600.00 (10)	529.80 (8)
44 (10)	600.00 (10)	600.00 (10)	470.61 (7)	600.00 (10)	600.00 (10)	553.19 (9)
47 (10)	600.00 (10)	600.00 (10)	485.60 (7)	600.00 (10)	600.00 (10)	529.00 (8)
50 (10)	600.00 (10)	600.00 (10)	485.07 (7)	600.00 (10)	600.00 (10)	526.66 (8)

Table 3: RSPAPER experiments: variable number of papers/reviewers, multiple answer sets, randomized.

experiments with instances that had few external atoms for eliminating answer set candidates but many local constraints. For such highly constrained instances, the new decomposition framework is not beneficial as it incurs an overhead compared to the monolithic evaluation that increases runtimes.

6.3.3 Summary

The results demonstrate a clear improvement using the new framework; they can often be further improved by fine-tuning the grounding and solving algorithm, and by customizing the default heuristics of the framework, as discussed by 16 (16; 15), and 46). However, already the default settings yield results that are significantly better than using the previous evaluation method or using no framework at all; note that the latter requires an advanced grounding algorithm as by 15), which was not available at the time the initial evaluation approach as by 48) was developed.

In conclusion, the evaluation framework in Section 5 pushes HEX-programs towards scalability for realistic instance sizes, which previous evaluation techniques missed.

7 Related Work and Discussion

We now discuss our results in the context of related work, and will address possible optimizations.

7.1 Related Work

7.1.1 External Sources

The *dlv-ex* system (6) was a pioneering work on value invention through external atoms in ASP. It supported VI-restricted programs, which amount to HEX-programs under extensional semantics without higher-order atoms and a strong safety condition that is subsumed by *lde-safety*. Answer set computation followed the traditional approach on top of *dlv*, but used a special progressive grounding method (thus an experimental comparison to solving, i.e., model building as in the focus of this paper, is inappropriate).

Also the *clingcon* system (40) is related, which enhances ASP with constraint atoms; this can be seen as a special case of HEX-programs that focuses on a particular external source. As for evaluation, an important difference to general external sources is that constraint atoms do not use value invention. The modularity techniques from above are less relevant for this setting as grounding the overall program in one shot is possible. However, this also fits into our framework as disabling decomposition in fact corresponds to a dedicated (trivial) heuristics which keeps the whole program as a single unit.

We also remark that *gringo* and *clasp* use a concept called “external atoms” for realizing various applications such as constraint ASP solving as in *clingcon* and incremental solving (23). However, despite their name they are different from external atoms in HEX-programs. In the former case, external atoms are excluded from grounding-time optimization such that these atoms are not eliminated even if their truth value is deterministically false during grounding. This allows to add rules that found truth of such atoms in later incremental grounding steps. In case of HEX the truth value is determined by external sources. Moreover *gringo* contains an interface for Lua and Python functions that can perform computations during grounding. HEX external atoms are more expressive: they cannot be evaluated during grounding because their semantics is defined with respect to the answer set.

7.1.2 Rule Dependencies

In the context of answer set programming, dependency graphs over rules have been used earlier, e.g. by (35) and (36). However, these works consider only ordinary ground programs, and furthermore the graphs are used for characterizing and computing the answer sets of a program from these graphs. In contrast, we consider nonground programs with and external atoms, and we use the graph to split the program into evaluation units with the goal of modularly computing answer sets.

7.1.3 Modularity

Our work is naturally related to work on program modularity under stable model semantics, as targeted by splitting sets (34) and descendants, with the work by (39) and (30) a prominent representative that lifted them to modular programs with choice rules and disjunctive rules, by considering “symmetric splitting”. Other works, e.g. by (33) go further to *define* semantics of systems of program modules, departing from a mere semantics-preserving decomposition of a larger program into smaller parts, or consider multi-language systems that combine modules in possibly different formalisms on equal terms (cf. e.g. (31) and (52)).

Comparing the works by (39) and (30) as, from a semantic decomposition perspective, the closest in this group to ours, an important difference is that our approach works for non-ground programs and explicitly considers possible overlaps of modules. It is tailored to efficient evaluation of arbitrary programs, rather than to facilitate module-style logic programming with declarative specifications, or to provide compositional semantics for modules beyond uni-lateral evaluation, as done by (31) and (52); for them, introducing values outside the module domain (known as *value invention*) does not play a visible role. In this regard, it is in line with previous HEX-program evaluation (20) and decomposition techniques to ground ordinary programs efficiently (7).

7.1.4 Splitting Theorems

Our new splitting theorems compare to related splitting theorems as follows.

Theorem 1 is similar to Theorem 4.6.2 by (48); however, we do not use splitting sets on atoms, but splitting sets on rules. Furthermore, (48) has no analog to Theorem 2.

The seminal Splitting Set Theorem by 34) divides the interpretation of P into disjoint sets X and Y , where X is an answer set of the ‘bottom’ $gb_A(P) \subseteq P$ and Y is an answer set of a ‘residual’ program obtained from $P \setminus gb_A(P)$ and X . In the residual program, all references to atoms in X are removed, in a way that it semantically behaves as if facts X were added to $P \setminus gb_A(P)$, while the answer sets of the residual do not contain any atom in X . This works nicely for answer set programs, but it is problematic when applied to HEX programs, because external atoms may depend on the bottom and on atoms in heads of the residual program; hence, they cannot be eliminated from rule bodies. The only way to eliminate bottom facts from the residual program would be to “split” external atoms semantically into a part depending on the bottom and the program remainder, and by replacing external atoms in rules with external atoms that have been partially evaluated wrt. a bottom answer set. Technically, this requires to introduce new external atoms, and formulating a splitting theorem for HEX programs with two disjoint interpretations X and Y is not straightforward. Furthermore, such external atom splitting and partial evaluation might not be possible in a concrete application scenario.

Different from the two splitting theorems recalled above, the Global Splitting Theorem by 20) does not split an interpretation of the program P into disjoint interpretations X and Y , and thus should be compared to our Theorem 2. However, the Global Splitting Theorem does not allow constraint sharing, and it involves a residual program which specifies how external atoms are evaluated via ‘replacement atoms’, which lead to extra facts D in the residual program that must be removed from its answer sets. Both the specification of replacement atoms and the extra facts make the Global Splitting Theorem cumbersome to work with when proving correctness of HEX encodings. Moreover, the replacement atoms are geared towards a certain implementation technique which however is not mandatory and can be avoided.

Lemma 5.1 by 17) is structurally similar to our Theorem 2: answer sets of the bottom program are evaluated together with the program depending on the bottom (here called the residual), hence answer sets of the residual are answer sets of the original program. However, the result was based on atom dependencies and did neither consider negation nor external atoms.

7.2 Possible Optimizations

Evaluation graphs naturally encode parallel evaluation plans. We have not yet investigated the potential benefits of this feature in practice, but this property allows us to do parallel solving based on solver software that does not have parallel computing capabilities itself (‘parallelize from outside’). This applies both to programs with external atoms, as well as to ordinary ASP programs (i.w., without external atoms). Improving reasoning performance by decomposition has been investigated by 1), however, only wrt. monotonic logics.

Improving HEX evaluation efficiency by using knowledge about domain restrictions of external atoms has been discussed by 13). These rewriting methods yield partially grounded sets of rules which can easily be distributed into distinct evaluation units by an optimizer. This directly provides efficiency gains as described in the above work.

As a last remark on possible optimizations, we observe that the data flow between evaluation units can be optimized using proper notions of model projection, such as in (24). Model projections would tailor input data of evaluation units to necessary parts of intermediate answer sets; however, given that different units might need different parts of the same intermediate input answer set, a space-saving efficient projection technique is not straightforward.

8 Conclusion

HEX-programs extend answer set programs with access to external sources through an API-style interface, which has been fruitfully deployed to various applications. Providing efficient evaluation methods for such programs is a challenging but important endeavor, in order to enhance the practicality of the approach and to make it eligible for a broader range of applications. In this direction, we have presented in this article a novel evaluation method for HEX-programs based on modular decomposition. We have presented new results for the latter using special splitting sets, which are more general than previous results and use rule sets as a basis for splitting rather than sets of atoms as in previous approaches. Furthermore, we have presented an evaluation framework which employs besides a traditional evaluation graph that consists of program components and reflects syntactic dependencies among them, also a model graph whose nodes collect answer sets that are combined and passed on between components. Using decomposition techniques, evaluation units can be dynamically formed and evaluated in the framework using different heuristics. Moreover, the answer sets of the overall program can be produced in a streaming fashion. The new approach leads in combination with other techniques to significant improvements for a variety of applications, as demonstrated by 15 (15; 16) and 46). Notably, while our results target HEX-programs, the underlying concepts and techniques are not limited to them (e.g., to separate the evaluation and the model graph) and may be fruitfully transferred to other rule-based formalisms.

8.1 Outlook

The work we presented can be continued in different directions. As for the prototype reasoner, a rather straightforward extension is to support brave and cautious reasoning on top of HEX programs, while incorporating constructs like aggregates or preference constraints requires more care and efforts. Regarding program evaluation, our general evaluation framework provides a basis for further optimizations and evaluation strategies. Indeed, the generic notions of evaluation unit, evaluation graph and model graph allow to specialize and improve our framework in different respects. First, evaluation units (which may contain duplicated constraints), can be chosen according to a proper estimate of the number of answer sets (the fewer, the better); second, evaluation plans can be chosen by ad-hoc optimization modules, which may give preference to (a combination of) time, space, or parallelization conditions. Third, the framework is amenable to a form of coarse-grained distributed computation at the level of evaluation units (in the style of 42)).

While modular evaluation is advantageous in many applications, it can also be counterproductive, as currently the propagation of knowledge learned by conflict-driven techniques into different evaluation units is not possible. In such cases, evaluating the program as a single evaluation unit is often also infeasible due to the properties of the grounding algorithm, as observed by 15). Thus, another starting point for future work is a tighter integration of the solver instances used to evaluate different units, e.g., by exchanging learned knowledge. In this context, also the interplay of the grounder and the solver is an important topic.

A Proofs

Proof of Theorem 1 (Splitting Theorem). Given a set of ground atoms M and a set of rules R , we denote by $M|_R = M \cap gh(R)$ the projection of M to ground heads of rules in R .

(\Rightarrow) Let $M \in \mathcal{AS}(P)$. We show that (1) $M|_R \in \mathcal{AS}(R)$ and that (2) $M \in \mathcal{AS}(P \setminus R \cup facts(M|_R))$.

As for (1), we first show that $M|_R$ satisfies the reduct $fR^{M|_R}$, and then that $M|_R$ is indeed a minimal model of $fR^{M|_R}$. M satisfies fP^M and $R \subseteq P$. Observe that, by definition of FLP reduct, $fR^M \subseteq fP^M$.

By definition of rule splitting set, satisfiability of rules in R does not depend on heads of rules in $P \setminus R$ (due to the restriction of external atoms to extensional semantics, this is in particular true for external atoms in R). Therefore $fR^{M|_R} = fR^M$, M satisfies $fR^{M|_R}$, and $M|_R$ satisfies $fR^{M|_R}$. For showing $M|_R \in \mathcal{AS}(R)$, it remains to show that $M|_R$ is a minimal model of $fR^{M|_R}$.

Assume towards a contradiction that some $S \subset M|_R$ is a model of $fR^{M|_R}$. Then there is a nonempty set $A = M|_R \setminus S$ of atoms with $A \subseteq gh(R)$. Let $M^* = M \setminus A$. We next show that M^* is a model of fP^M , which implies that $M \notin \mathcal{AS}(P)$. Assume on the contrary that M^* is not a model of fP^M . Hence there exists some rule $r \in fP^M$ such that $H(r) \cap M^* = \emptyset$, $B^+(r) \subseteq M^*$, $B^-(r) \cap M^* = \emptyset$ and external atoms in $B^+(r)$ (resp., $B^-(r)$) evaluate to true (resp., false) wrt. M^* . S agrees with M^* on atoms from $gh(R)$, and S satisfies $fR^{M|_R}$. The truth values of external atoms in bodies of rules in R depends only on atoms from $gh(R)$, therefore external atoms in R evaluate to the same truth value wrt. S and M^* . Therefore $r \notin fR^{M|_R}$ and $r \in f(P \setminus R)^M$. Since $r \in P \setminus R$, $H(r) \subseteq gh(P \setminus R)$, and because M and M^* agree on atoms from $gh(P \setminus R)$, $H(r) \cap M^* = \emptyset$ from above implies that $H(r) \cap M = \emptyset$. Because $r \in fP^M$, its body is satisfied in M , and since its head has no intersection with M , we get that fP^M is not satisfied by M , which is a contradiction. Therefore M^* is a model of fP^M . As $M^* \subset M$, this contradicts our assumption that $M \in \mathcal{AS}(P)$. Therefore $S = M|_R = X$ is a minimal model of fR^M .

We next show that M satisfies the reduct $f(P \setminus R \cup facts(M|_R))^M$, and then that it is indeed a minimal model of the reduct. By the definition of reduct, $f(P \setminus R \cup facts(M|_R))^M = f(P \setminus R)^M \cup facts(M|_R)$. M satisfies $facts(M|_R)$ because $M|_R \subseteq M$. Furthermore $f(P \setminus R)^M \subseteq fP^M$, hence M satisfies $f(P \setminus R)^M$. Therefore M satisfies $f(P \setminus R \cup facts(M|_R))^M$.

To show that M is a minimal model of $f(P \setminus R \cup facts(M|_R))^M$, assume towards a contradiction that some $S \subset M$ is a model of $f(P \setminus R \cup facts(M|_R))^M$. Since $facts(M|_R)$ is part of the reduct, $M|_R \subseteq S$, therefore $S|_{gh(R)} = M|_R$. By definition of rule splitting set, satisfiability of rules in R does not depend on heads of rules in $P \setminus R$, hence S satisfies fR^M . Because S satisfies $f(P \setminus R \cup facts(M|_R))^M = f(P \setminus R)^M \cup facts(M|_R)$, it also satisfies $f(P \setminus R)^M$. Since S satisfies both fR^M , S satisfies $fP^M = f(P \setminus R)^M \cup fR^M$. This is a contradiction to $M \in \mathcal{AS}(P)$. Therefore $S = M$ is a minimal model of $f(P \setminus R \cup facts(M|_R))^M$.

(\Leftarrow) Let $M \in \mathcal{AS}(P \setminus R \cup facts(X))$ and let $X \in \mathcal{AS}(R)$. We first show that M satisfies fP^M , and then that it is a minimal model of fP^M .

As facts X are part of the program $P \setminus R \cup facts(X)$, and by definition of rule splitting set, $P \setminus R$ contains no rule heads unifying with $gh(R)$, hence we have $X = M|_R$. Furthermore $f(P \setminus R \cup facts(X))^M \setminus facts(X) \cup fR^M = fP^M$, and as M satisfies the left side, it satisfies the right side. To show that M is a minimal model of fP^M , assume $S \subset M$ is a smaller model of fP^M . By definition of reduct, S also satisfies $f(P \setminus R)^M$ and fR^M . Since R is a splitting set, satisfiability of rules in R does not depend on heads of rules in $P \setminus R$, therefore $fR^M = fR^{M|_R} = fR^X$ and $S|_{gh(R)}$ satisfies fR^X . Since $S \subset M$, we have $S|_{gh(R)} \subseteq X$. Because X is a minimal model of fR^X , $S|_{gh(R)} \subset X$ is impossible and $S|_{gh(R)} = X$. Therefore $S|_{gh(P \setminus R)} \subset M|_{gh(P \setminus R)}$. Because S satisfies $f(P \setminus R)^M$ and $S|_{gh(R)} = X$, S also satisfies $f(P \setminus R \cup facts(X))^M$. Since $S \subset M$, this contradicts the fact that M is a minimal model of $P \setminus R \cup facts(X)$. Therefore $S = M$ is a minimal model of fP^M . \square

Proof of Theorem 2 (Generalized Splitting Theorem). By definition of generalized bottom, the set $C = B \setminus R$ contains only constraints, therefore $gh(B) = gh(R)$ and $M|_{gh(B)} = M|_{gh(R)}$. As $R \subseteq B$ and $B \setminus R$ contains only constraints, $\mathcal{AS}(B) \subseteq \mathcal{AS}(R)$. The only difference between Theorem 1 and Theorem 2 is, that for obtaining X , the latter takes additional constraints into account.

(\Rightarrow) It is sufficient to show that $M|_{gh(B)}$ does not satisfy the body of any constraint in $C \subseteq P$ if M does not satisfy the body of any constraint in P . Since B is a generalized bottom, no negative dependencies of

constraints C to rules in $P \setminus B$ exist; therefore if the body of a constraint $c \in C$ is not satisfied by M , the body of c is not satisfied by $M|_{gh(B)}$. As M satisfies P , it does not satisfy any constraint body in P , hence the projection $M|_{gh(B)}$ does not satisfy any constraint body in $B \setminus R$.

(\Leftarrow) It is sufficient to show that an answer set of R that satisfies a constraint body in C also satisfies that constraint body in P , which raises a contradiction. As constraints in C have no negative dependencies to rules in $P \setminus B$, a constraint with a satisfied body in $M|_{gh(R)}$ also has a satisfied body in M , therefore the result follows. \square

Proof of Proposition 1. Assume towards a contradiction that there exist a non-constraint $r \in P$, a rule $s \in P$ with $r \rightarrow_{m,n} s$, and $u' \in U|_r, v' \in U|_s$ such that $(u', v') \notin E$. Due to Definition 9, $r \rightarrow_{m,n} s$ implies that s has $H(s) \neq \emptyset$ and therefore that s is a non-constraint. Definition 14 (b) then implies that $U|_r = \{u'\}$ and $U|_s = \{v'\}$ (non-constraints are present in exactly one unit).

Case (i): for $r \rightarrow_n s$, Definition 14 (c) specifies that for all $u \in U|_r$ and $v \in U|_s$ there exists an edge $(u, v) \in E$, therefore also $(u', v') \in E$, which is a contradiction.

Case (ii): for $r \rightarrow_m s$, Definition 14 (d) specifies that some $u \in U|_r$ exists such that for every $v \in U|_s$ there exists an edge $(u, v) \in E$; since $U|_r = \{u'\}$ and $U|_s = \{v'\}$, it must hold that $(u', v') \in E$, which is a contradiction. \square

Proof of Proposition 2. Given two distinct units $u_1, u_2 \in U$, assume towards a contradiction that some $\gamma \in gh(u_1) \cap gh(u_2)$ exists. Then there exists some $r \in u_1$ with $\alpha \in H(r)$ and $\alpha \sim \gamma$, and there exists some $s \in u_2$ with $\beta \in H(s)$ and $\beta \sim \gamma$. As $\alpha \sim \gamma$ and $\beta \sim \gamma$ and γ is ground, we obtain $\alpha \sim \beta$; hence, by Definition 9 (iii) we have $r \rightarrow_m s$ and $s \rightarrow_m r$. As r and s have nonempty heads, they are non-constraints. Thus by Proposition 1, there exist edges $(u_1, u_2), (u_2, u_1) \in E$. As an evaluation graph is acyclic, it follows $u_1 = u_2$; this is a contradiction. \square

Proof of Proposition 3. For an Ide-safe program P , the graph $\mathcal{E} = (\{P\}, \emptyset)$ is a valid evaluation graph. \square

Proof of Theorem 3. For any set of rules, let $constr(S) = \{r \in S \mid B(r) = \emptyset\}$ denote the set of constraints in S . We say that the *dependencies of $r \in Q$ are covered at unit $u \in U$* , if for every rule $s \in Q$ such that $r \rightarrow_{m,n} s$ and $s \notin u$, it holds that $(u, u') \in E$ for all $u' \in U|_s$, i.e., u has an edge to all units containing s .

To prove that $B = u^<$ is a generalized bottom of $P = u^{\leq}$ wrt. the rule splitting set $R = u^< \setminus constr(u^<)$ as by Definition 12, we prove that (a) $R \subseteq B \subseteq P$, (b) $B \setminus R$ contains only constraints, (c) no constraint in $B \setminus R$ has nonmonotonic dependencies to rules in $P \setminus B$, and (d) R is a rule splitting set of P .

Statement (a) corresponds to $u^< \setminus constr(u^<) \subseteq u^< \subseteq u^{\leq}$ and u^{\leq} is defined as $u^{\leq} = u^< \cup u$, therefore the relations all hold. For (b), $B \setminus R = u^< \setminus (u^< \setminus constr(u^<))$, and as $A \setminus (A \setminus B) = A \cap B$, it is easy to see that $B \setminus R = u^< \cap constr(u^<)$ and thus $B \setminus R$ only contains constraints. For (c), we show a stronger property, namely that no rule (constraint or non-constraint) in B has nonmonotonic dependencies to rules in $P \setminus B$. $B = u^<$ is the union of evaluation units $V = \{v \in U \mid v < u\}$. By Definition 14 (c) all nonmonotonic dependencies $r \rightarrow_n s$ are covered at every unit w such that $w \in U_r$. Hence if $r \in w$ and $w \in V$, then either $s \in w$ or $s \in w^<$ holds, and hence $s \in w^{\leq} \subseteq u^<$. As $P \setminus B = u^{\leq} \setminus u^<$, no nonmonotonic dependencies from $B = u^<$ to $P \setminus B$ exist and (c) holds. For (d) we know that $R = u^< \setminus constr(u^<)$ contains no constraints, and by Proposition 1 all dependencies of non-constraints in R are covered by \mathcal{E} . Therefore $r \in R, r \rightarrow_{m,n} s$, and $s \in P$ implies that $s \in R$. Consequently, (d) holds which proves the theorem. \square

Proof of Theorem 4. Similar to the proof of Theorem 3, we show this in four steps; given $P = u^<$, $R = u'^{\leq} \setminus \text{constr}(u'^{\leq})$, and $B = u'^{\leq} = u' \cup u'^{<}$, we show that (a) $R \subseteq B \subseteq P$, (b) $B \setminus R$ contains only constraints, (c) no constraint in $B \setminus R$ has nonmonotonic dependencies to rules in $P \setminus B$, and (d) R is a rule splitting set of P . Let $\text{preds}_{\mathcal{E}}(u) = \{u_1, \dots, u_k\}$ and Let $V = \{v \in U \mid v < u'\}$ be the set of units on which u' transitively depends. (Note that $V \subset \text{preds}_{\mathcal{E}}(u)$ and $u \notin V$.) As $u'^{<}$ contains all units u' transitively depends on, we have $B = u' \cup \bigcup_{w \in V} w$.

For (a), $R \subseteq B$ holds trivially, and $B \subseteq P$ holds by definition of $u^<$ and u'^{\leq} and because $u' \in \text{preds}_{\mathcal{E}}(u)$. Statement (b) holds, because $B \setminus R$ removes R from B , i.e., it removes everything that is not a constraint in B from B , therefore only constraints remain. For (c) we show that no rule in B has a nonmonotonic dependency to rules in $P \setminus B$. By Definition 14 (c), all nonmonotonic dependencies are covered at all units. Therefore a rule $r \in w$, $w \in \{u'\} \cup V$ with $r \rightarrow_n s$, $s \in U$ implies that either $s \in w$, or that s is contained in a predecessor unit of w and therefore in u' or in V . Hence there are no nonmonotonic dependencies from rules in B to any rules not in B , and hence also not to rules in $P \setminus B$ and (c) holds. For (d) we know that R contains no constraints and by Proposition 1 all dependencies of non-constraints in R are covered by \mathcal{E} . Therefore $r \in R$, $r \rightarrow_{m,n} s$, $s \in P$ implies that $s \in R$ and the theorem holds. \square

Proof of Proposition 4. (\Rightarrow) The added vertex m' is assigned to one unit and gets assigned a type. Furthermore, the graph stays acyclic as only outgoing edges from m' are added. I-connectedness is satisfied, as it is satisfied in \mathcal{I} and we add no o-interpretation. O-connectedness is satisfied, as m' gets appropriate edges to o-interpretations at its predecessor units, and for other i-interpretations it is already satisfied in \mathcal{I} .

For FAI intersection, observe that if we add an edge (m', m_i) to \mathcal{I} and it holds that $m_i \in o\text{-ints}_{\mathcal{I}}(u_i)$, then m' reaches in \mathcal{I} only one o-interpretation at u_i , and due to O-connectedness that o-interpretation is connected to exactly one i-interpretation at u_i , which is part of the original graph \mathcal{I} and therefore satisfies FAI intersection. Therefore it remains to show that the union of subgraphs of \mathcal{I} reachable in \mathcal{I} from m_1, \dots, m_k , contains one o-interpretation at each unit in the subgraph of \mathcal{E} reachable from u_1, \dots, u_k . We make a case distinction.

Case (I): two o-interpretations $m_i \in o\text{-ints}_{\mathcal{I}}(u_i)$, $m_j \in o\text{-ints}_{\mathcal{I}}(u_j)$ in the join, with $1 \leq i < j \leq k$, have no common unit that is reachable in \mathcal{E} from u_i and from u_j : then the condition is trivially satisfied, as the subgraphs of \mathcal{I} reachable in \mathcal{I} from m_i and m_j , respectively, do not intersect at any unit.

Case (II): two o-interpretations $m_i \in o\text{-ints}_{\mathcal{I}}(u_i)$, $m_j \in o\text{-ints}_{\mathcal{I}}(u_j)$ in the join, with $1 \leq i < j \leq k$, have at least one common unit that is reachable from u_i and from u_j in \mathcal{E} . Let u^f be a unit reachable in \mathcal{E} from both u_i and u_j on two paths that do not intersect before reaching u^f . From u_i to u^f , and from u_j to u^f , exactly one o-interpretation is reachable in \mathcal{I} from m_i and m_j , respectively, as these paths do not intersect. u^f is a FAI of u , and as the join is defined, we reach in \mathcal{E} exactly one o-interpretation at unit u^f from m_i and m_j . Due to O-connectedness, we also reach in \mathcal{I} exactly one i-interpretation m'' at u^f from m_i and m_j . Now m'' is common to subgraphs of \mathcal{I} that are reachable in \mathcal{I} from m_i and m_j , and m'' satisfies FAI intersection in \mathcal{I} .

Consequently, FAI intersection is satisfied in \mathcal{I}' for all pairs of predecessors of m' and therefore in all cases. As no vertex m with $\{(m, m_1), \dots, (m, m_k)\} \subseteq F$ exists and as \mathcal{I} satisfies Uniqueness, also \mathcal{I}' satisfies Uniqueness.

(\Leftarrow) Assume towards a contradiction that \mathcal{I}' is an i-graph but that the join is not defined. Then there exists some FAI $u' \in \text{fai}(u)$ such that either no or more than one o-interpretation from $o\text{-ints}_{\mathcal{I}}(u)$ is reachable in \mathcal{I} from some m_i , $1 \leq i \leq k$. As \mathcal{I} is an i-graph, due to I-connectedness and O-connectedness, if a unit u' is a FAI and therefore u' is reachable in \mathcal{E} from u_i , then at least one i-interpretation and one o-interpretation at u' is reachable in \mathcal{I} from m_i . If more than one o-interpretation is reachable in \mathcal{I} from

some m_i , $1 \leq i \leq k$, this means that more than one o-interpretation at u' is reachable in \mathcal{T}' from the newly added i-interpretation m . However, this violates FAI intersection in \mathcal{T}' , which is a contradiction. Hence the result follows. \square

Proof of Proposition 5. (\Rightarrow) Whenever the join is defined, \mathcal{A}' is an i-graph by Proposition 4. It remains to show that $\text{int}(m')^+ \in \mathcal{AS}(u^<)$, and that \mathcal{A}' fulfills items (a) and (c) of an answer set graph. By Theorem 4 we know that for each u_i , $u_i^<$ is a generalized bottom of $u^<$ wrt. the set $R_i = \{r \in u_i^< \mid B(r) \neq \emptyset\}$. For each u_i , therefore $Y \in \mathcal{AS}(u^<)$ iff $Y \in \mathcal{AS}(u^< \setminus R_i \cup \text{facts}(X))$ for some $X \in \mathcal{AS}(u_i^<)$. As \mathcal{A} is an answer set graph, for each m_i we know that $\text{int}(m_i)^+ \in \mathcal{AS}(u_i^<)$; hence $Y \in \mathcal{AS}(u^<)$ if $Y \in \mathcal{AS}(u^< \setminus R_i \cup \text{int}(m_i)^+)$. Now from the evaluation graph properties we know that $u^< = u_1^< \cup \dots \cup u_k^<$, and from the construction of $\text{int}(m')$ and its dependencies in \mathcal{A}' we obtain that $\text{int}(m')^+ = \text{int}(m_1)^+ \cup \dots \cup \text{int}(m_k)^+$. It follows that $\text{int}(m')^+ \in \mathcal{AS}(u^<)$, which satisfies condition (a). Due to the definition of join, condition (c) is also satisfied and \mathcal{A}' is indeed an answer set graph.

(\Leftarrow) As \mathcal{A}' is an answer set graph, it is an i-graph, and hence by Proposition 4 $m = m_1 \bowtie \dots \bowtie m_k$ is defined. \square

Proof of Theorem 5. We prove this theorem using Proposition 6. We construct $\mathcal{E}'' = (U'', E'')$ with $U'' = U \cup \{u_{\text{final}}\}$, $u_{\text{final}} = \emptyset$, and $E'' = E \cup \{(u_{\text{final}}, u) \mid u \in U\}$. As u_{final} contains no rules and as \mathcal{E}'' is acyclic, no evaluation graph property of gets violated and \mathcal{E}'' is also an evaluation graph. As \mathcal{A} contains no interpretations at u_{final} and dependencies from units in U are the same in \mathcal{E} and \mathcal{E}'' , \mathcal{A} is in fact an answer set graph for \mathcal{E}'' . We now modify \mathcal{A} to obtain \mathcal{A}'' as follows. We add the set $M_{\text{new}} = \{m \mid m = m_1 \bowtie \dots \bowtie m_n \text{ is defined at } u_{\text{final}} \text{ (wrt. } \mathcal{A})\}$ as i-interpretations of u_{final} and dependencies from each $m \in M_{\text{new}}$ to the respective o-interpretations m_i , $1 \leq i \leq n$. By Proposition 5, \mathcal{A}'' is an answer set graph for \mathcal{E}'' , and moreover \mathcal{A}'' gets input-complete for u_{final} by construction. As \mathcal{A}'' is input-complete for $U \cup \{u_{\text{final}}\}$ and output-complete for U , by Proposition 6 we have that $\mathcal{AS}(P) = i\text{-ints}_{\mathcal{A}}(u_{\text{final}}) = M_{\text{new}}$. As for every join $m = m_1 \bowtie \dots \bowtie m_n$, we have $\text{int}(m) = \text{int}(m_1) \cup \dots \cup \text{int}(m_n)$, to complete the proof of the theorem, it remains to show that the join m between m_1, \dots, m_n is defined at u_{final} iff the subgraph \mathcal{A}' of \mathcal{A} reachable from the o-interpretations m_i in F fulfills $|o\text{-ints}_{\mathcal{A}}(u_i)| = 1$, for each $u_i \in U$. As the join involves all units in U , and since \mathcal{A}'' is an answer set graph and thus an i-graph, it follows from the conditions for an i-graph that at each $u_i \in U$ exactly one o-interpretation is reachable from m , and thus also from each m_i ; thus the condition for \mathcal{A}' holds. Conversely, if the subgraph \mathcal{A}' fulfills $|o\text{-ints}_{\mathcal{A}}(u_i)| = 1$ for each $u_i \in U$, then clearly the FAI condition for the join m being defined is fulfilled. \square

Proof of Proposition 6. As u_{final} depends on all units in $U \setminus \{u_{\text{final}}\}$, due to O-connectedness every i-interpretation $m \in i\text{-ints}_{\mathcal{A}}(u_{\text{final}})$ depends on one o-interpretation at every unit in $U \setminus \{u_{\text{final}}\}$. Let $U \setminus \{u_{\text{final}}\} = \{u_1, \dots, u_k\}$ and let $M_M = \{m_1, \dots, m_k\}$ be the set of o-interpretations such that $(m, m_i) \in F$ and $m_i \in o\text{-ints}_{\mathcal{A}}(u_i)$, $1 \leq i \leq k$. Then, due to FAI intersection, M_m contains each o-interpretation that is reachable from m in \mathcal{A} , and M_m contains only interpretations with this property. Hence $\text{int}(m)^+ = \text{int}(m_1) \cup \dots \cup \text{int}(m_k)$, and due to condition (c) in Definition 19, we have $\text{int}(m) = \text{int}(m)^+$. By the dependencies of u_{final} , we have $u_{\text{final}}^< = P$, and as u_{final} is input-complete, we have that $\mathcal{AS}(P) = \mathcal{AS}(u_{\text{final}}^<) = \{\text{int}(m)^+ \mid m \in i\text{-ints}_{\mathcal{A}}(u_{\text{final}})\}$. As $\text{int}(m) = \text{int}(m)^+$ for every i-interpretation m at u_{final} , we obtain the result. \square

Proof of Proposition 7. The proposition follows from Property 1, which asserts that the grounding P' has the same answer sets as P , and from the soundness and completeness of the evaluation algorithm for ground HEX-programs as asserted by Property 2. \square

Proof of Theorem 6. We show by induction on its construction that $\mathcal{I} = (M, F, \text{unit}, \text{type}, \text{int})$ is an answer set graph for \mathcal{E} , and that at the beginning of the while-loop \mathcal{I} is input- and output-complete for $V \setminus U$.

(Base) Initially, \mathcal{I} is initially and $V = U$, hence the base case trivially holds.

(Step) Suppose that \mathcal{I} is an answer set graph for \mathcal{E} at the beginning of the while-loop, and that it is input- and output-complete for $V \setminus U$. As the chosen u only depends on units in $V \setminus U$, it depends only on output-complete units. For a leaf unit u , (b) creates an empty i-interpretation and therefore makes u input-complete. For a non-leaf unit u , the first for-loop (c) builds all possible joins of interpretations at predecessors of u and adds them as i-interpretations to \mathcal{I} . As all predecessors of u are output-complete by the hypothesis, this makes u input-complete. Now suppose that Condition (d) is false, i.e., $u \neq u_{\text{final}}$. Then the second for-loop (e) evaluates u wrt. every i-interpretation at u and adds the result to u as an o-interpretation. Due to Proposition 7, $\text{EVALUATELDES SAFE}(u, \text{int}(m'))$ returns all interpretations o such that $o \in \{X \setminus \text{int}(m') \mid X \in \mathcal{AS}(u \cup \text{facts}(\text{int}(m')))\}$. As u depends on all units on which its rules depend, and as i-interpretations contain all atoms from o-interpretations of predecessor units (due to condition (c) of Definition 19), we have $\text{EVALUATELDES SAFE}(u, \text{int}(m')) = \text{EVALUATELDES SAFE}(u, \text{int}(m')^+)$. By Theorem 3, $u^<$ is a generalized bottom of u^{\leq} , and by the induction hypothesis $\text{int}(m')^+ \in \mathcal{AS}(u^<)$; hence by Theorem 2, we have that $\text{int}(m')^+ \cup o \in \mathcal{AS}(u^{\leq})$. Consequently, adding a new o-interpretation m with interpretation $\text{int}(m) = o$ and dependency to m' to the graph \mathcal{I} results in $\text{int}(m)^+ \in \mathcal{AS}(u^{\leq})$, and adding all of them makes \mathcal{I} output-complete for u . Finally, in (f) u is removed from U ; hence at the end of the while-loop \mathcal{I} is an answer set graph and again input- and output-complete for $V \setminus U$.

It remains to consider the case where Condition (d) is true. Then u_{final} was made input-complete, which means that all predecessors of u_{final} are output-complete. As u_{final} depends on all other units, we have $U = \{u_{\text{final}}\}$ and the algorithm returns $i\text{-ints}_{\mathcal{A}}(u)$; by Proposition 6, it thus returns $\mathcal{AS}(P)$, which will happen in the $|V|$ -th iteration of the while loop. \square

B Example Run of Algorithm 2

We provide here an example run of Algorithm 2 for our running example.

Example 30 (ctd.) Consider an evaluation graph \mathcal{E}'_2 which is \mathcal{E}_2 plus $u_{\text{final}} = \emptyset$, which depends on all other units. Following Algorithm 2 we first choose $u = u_1$, and as u_1 has no predecessor units, step (b) creates the i-interpretation m_1 with $\text{int}(m_1) = \emptyset$. As $u_1 \neq u_{\text{final}}$, we continue and in loop (e) obtain $O = \mathcal{AS}(u_1) = \{\{\text{swim}(\text{in})\}, \{\text{swim}(\text{out})\}\}$. We add both answer sets as o-interpretations m_2 and m_3 and then finish the outer loop with $U = \{u_2, u_3, u_4, u_{\text{final}}\}$. In the next iteration, we could choose $u = u_2$ or $u = u_3$; assume we choose u_2 . Then $\text{preds}_{\mathcal{E}}(u_2) = \{u_1\}$ and $k = 1$, and we enter the loop (c) and build all joins that are possible with o-interpretations at u_1 (all joins are trivial and all are possible), i.e., we copy the interpretations and store them at u_2 as new i-interpretations m_4 and m_5 . In the loop (e), we obtain $O = \text{EVALUATELDES SAFE}(u_2, \{\text{swim}(\text{in})\}) = \emptyset$, as indoor swimming requires money which is excluded by $c_8 \in u_2$. Therefore i-interpretation $\{\text{swim}(\text{in})\}$ yields no o-interpretation, indicated by ζ . However, we obtain $O = \text{EVALUATELDES SAFE}(u_2, \{\text{swim}(\text{out})\}) = \{\emptyset\}$: as outdoor swimming neither requires money nor anything else, i-interpretation $\{\text{swim}(\text{out})\}$ derives no additional atoms and yields the empty answer set, which we store as o-interpretation m_6 at u_2 ; the iteration ends with $U = \{u_3, u_4, u_{\text{final}}\}$. In the next iteration we choose $u = u_3$, we add in loop (c) i-interpretations m_7 and m_8 to u_3 , and in loop (e) o-interpretations m_9, \dots, m_{12} to u_3 ; the iteration ends with $U = \{u_4, u_{\text{final}}\}$. In the next iteration we choose $u = u_4$; this time we have multiple predecessors, and in loop (c) we check join candidates $m_6 \bowtie m_9$

Algorithm 3: ANSWERSETSONDEMAND

Input: evaluation graph \mathcal{E} for program P , with final unit $u_{final} = \emptyset$
Output: the answer sets of P
initialize global storage \mathcal{S}
repeat
 $m_{out} := \text{GETNEXTOUTPUTMODEL}(u_{final})$
 if $m_{out} \neq \text{UNDEF}$ **then** output m_{out}
until $m_{out} = \text{UNDEF}$

and $m_6 \bowtie m_{10}$, which are both not defined. The other join candidates are $m_6 \bowtie m_{11}$ and $m_6 \bowtie m_{12}$, which are both defined; we thus add their results as i -interpretations m_{13} and m_{14} , respectively, to u_4 . The loop (e) computes then one o -interpretation m_{15} for i -interpretation m_{13} and no o -interpretation for m_{14} . The iteration ends with $U = \{u_{final}\}$. In the next iteration, we have $\text{preds}_{\mathcal{E}}(u_{final}) = \{u_1, u_2, u_3, u_4\}$ and the loop (c) checks all combinations of one o -interpretation at each unit in $\text{preds}_{\mathcal{E}}(u_{final})$. Only one such join candidate is defined, namely $m = m_3 \bowtie m_6 \bowtie m_{11} \bowtie m_{15}$, whose result is stored as a new i -interpretation at u_{final} . The check (d) now succeeds, and we return all i -interpretations at u_{final} ; i.e., we return $\{m\} = \{\{swim(out), goto(altD), ngoto(gansD), go, need(loc, yogamat)\}\}$. This is indeed the set of answer sets of P_{swim} . \square

C On Demand Model Streaming Algorithm

Algorithm 2 fully evaluates all other units before computing results at the final evaluation unit u_{final} , and it keeps the intermediate results in memory. If we are only interested in one or a few answer sets, many unused results may be calculated.

Using the same evaluation graph, we can compute the answer sets with a different, more involved algorithm ANSWERSETSONDEMAND (shown in Algorithm 3) that operates demand-driven from units, starting with u_{final} , rather than data-driven from completed units. It uses in turn several building blocks that are shown in Algorithms 4–6

ANSWERSETSONDEMAND calls Algorithm GETNEXTOUTPUTMODEL for u_{final} and outputs its output models, i.e., the answer sets of the input program P given by the evaluation graph \mathcal{E} , one by one until it gets back UNDEF. Like Algorithm 2, GETNEXTOUTPUTMODEL builds in combination with the other algorithms an answer set graph \mathcal{A} for \mathcal{E} that is input-complete at all units, if all statements marked with ‘(+)’ are included; omitting them, it builds \mathcal{A} virtually and has at any time at most one input and one output model of each unit in memory.

Roughly speaking, the models at units are determined in the same order in which a right-to-left depth-first-traversal of the evaluation graph \mathcal{E} would backtrack from edges. This is because first all models of the subgraph reachable from a unit u are determined, then models at the unit u , and then the algorithm backtracks. The models of the subgraph are retrieved with GETNEXTINPUTMODEL one by one, and using *NextAnswerSet* the output models are generated and returned. The latter function is assumed to return, given a HEX-program P and the i -th element in an arbitrary but fixed enumeration I_1, I_2, \dots, I_m of the answer sets of P (without duplicates), the next answer set I_{i+1} , where by convention $I_0 = \text{UNDEF}$ and the return value for I_m is UNDEF. This is easy to provide on top of current solvers, and the incremental usage of *NextAnswerSet* allows for an efficient stateful realization (e.g. answer set computation is suspended).

The trickiest part of this approach is GETNEXTINPUTMODEL, which has to create locally and in an incremental fashion all joins that are globally defined, i.e., all combinations of incrementally available output models of predecessors which share a common predecessor model at all FAIs. To generate all combinations

Algorithm 4: GETNEXTOUTPUTMODEL(u)

```

Input:  $u$ : unit
Output:  $m_{out}$ : next omodel at  $u$  or UNDEF
if  $refsO(u) > 0$  then return UNDEF
if  $curI(u) = \text{UNDEF}$  then  $curI(u) := \text{GETNEXTINPUTMODEL}(u)$ 
while  $curI(u) \neq \text{UNDEF}$  do
   $curO(u) := \text{NextAnswerSet}(u \cup \text{facts}(curI(u)), curO(u))$ 
  if  $curO(u) \neq \text{UNDEF}$  then
    (+)  $\left[ \begin{array}{l} \text{add omodel } curO(u) \text{ to } \mathcal{A} \text{ with dependency to } curI(u) \\ \text{return } curO(u) \end{array} \right.$ 
   $curI(u) := \text{GETNEXTINPUTMODEL}(u)$ 
return UNDEF

```

Algorithm 5: ENSUREMODELINCREMENT(u, at)

```

Input:  $u$ : unit with  $\{u_1, \dots, u_k\} = \text{preds}_{\mathcal{E}}(u)$ ,  $at$ : index  $1 \leq at \leq k$ 
Output:  $at'$ : index  $at \leq at' \leq k$  or UNDEF
repeat
   $refsO(u_{at}) := refsO(u_{at}) - 1$ 
   $m := \text{GETNEXTOUTPUTMODEL}(u_{at})$ 
  if  $m = \text{UNDEF}$  then  $at := at + 1$ 
  else
     $\left[ \begin{array}{l} refsO(u_{at}) := refsO(u_{at}) + 1 \\ \text{return } at \end{array} \right.$ 
until  $at = k + 1$ 
return UNDEF

```

of output models in the right order, it uses the algorithm ENSUREMODELINCREMENT.

The algorithms operate on a global data structure $\mathcal{S} = (\mathcal{E}, \mathcal{A}, curI, curO, refsO)$ called *storage*, where

- $\mathcal{E} = (U, E)$ is the evaluation graph containing $u_{final} \in U$,
- $\mathcal{A} = (M, F, unit, type, int)$ is the (virtually built) answer set graph,
- $curI : U \rightarrow M \cup \{\text{UNDEF}\}$ and $curO : U \rightarrow M \cup \{\text{UNDEF}\}$, are functions that informally associate with a unit u the current input respectively output model considered, and
- $refsO : U \rightarrow \mathbb{N} \cup \{0\}$ is a function that keeps track of how many current input models point to the current output model of u ; this is used to ensure correct joins, by checking in GETNEXTOUTPUTMODEL that the condition (IG-F) for sharing models in the interpretation graph is not violated (for details see Section 5.1.2 and Definition 17).

Initially, the storage \mathcal{S} is empty, i.e., it contains the input evaluation graph \mathcal{E} , an empty answer set graph \mathcal{A} , and the functions are set to $curI(u) = \text{UNDEF}$, $curO(u) = \text{UNDEF}$, and $refsO(u) = 0$ for all $u \in U$. The call of GETNEXTOUTPUTMODEL for u_{final} triggers the right-to-left depth-first traversal of the evaluation graph.

We omit tracing Algorithm ANSWERSETSONDEMAND on our running example, as this would take quite some space; however, one can check that given the evaluation graph \mathcal{E}_2 , it correctly outputs the single answer set

$$I = \{swim(out), goto(altD), ngoto(gansD), go, need(loc, yogamat)\}.$$

Formally, it can be shown that given an evaluation graph $\mathcal{E} = (U, E)$ of a program P such that \mathcal{E} contains a final unit $u_{final} = \emptyset$, Algorithm ANSWERSETSONDEMAND outputs one by one all answer sets of P , without duplicates, and that in the version without (+)-lines, it stores at most one input and one output model per unit (hence the size of the used storage is linear in the size of the ground program $grnd(P)$).

Algorithm 6: GETNEXTINPUTMODEL(u)

```

Input:  $u$ : unit
Output:  $m_{out}$ : imodel at  $u$  or UNDEF

(a) if  $preds_{\mathcal{E}}(u) = \emptyset$  then
    if  $curI(u) = \text{UNDEF}$  then
        (+) add imodel  $\emptyset$  at  $u$  to  $\mathcal{A}$ 
        return  $\emptyset$ 
    else return UNDEF

let  $\{u_1, \dots, u_k\} = preds_{\mathcal{E}}(u)$  /* assume this order is fixed for each unit  $u$  */
if  $curI(u) \neq \text{UNDEF}$  then
     $at := \text{ENSUREMODELINCREMENT}(u, 1)$ 
    if  $at = \text{UNDEF}$  then return UNDEF
     $at := at - 1$ 
else  $at := k$ 
(b) while  $at \neq 0$  do
    if  $curO(u_{at}) \neq \text{UNDEF}$  then
         $refsO(u_{at}) := refsO(u_{at}) + 1$ 
         $at := at - 1$ 
    else
         $m := \text{GETNEXTOUTPUTMODEL}(u_{at})$ 
        if  $m = \text{UNDEF}$  then
            if  $at = k$  then return UNDEF
             $at := \text{ENSUREMODELINCREMENT}(u, at + 1)$ 
            if  $at = \text{UNDEF}$  then return UNDEF
        else
             $refsO(u_{at}) := refsO(u_{at}) + 1$ 
             $at := at - 1$ 

let  $m = curO(u_1) \bowtie \dots \bowtie curO(u_k)$ 
(+ add imodel  $m$  to  $\mathcal{A}$  with dependencies to  $curO(u_1), \dots, curO(u_k)$ 
return  $m$ 

```

D Overview of Liberal Domain-Expansion Safety

Strong domain-expansion safety is overly restrictive, as it also excludes programs that clearly *are* finitely restrictable. In this section we give an overview about the notion and refer to (15) for details.

Example 31 Consider the following program:

$$P = \left\{ \begin{array}{l} r_1 : p(a). \quad r_3 : s(Y) \leftarrow p(X), \&concat[X, a](Y). \\ r_2 : q(aa). \quad r_4 : p(X) \leftarrow s(X), q(X). \end{array} \right\}$$

It is not strongly safe because Y in the cyclic external atom $\&concat[X, a](Y)$ in r_3 does not occur in an ordinary body atom that does not depend on $\&concat[X, a](Y)$. However, P is finitely restrictable as the cycle is “broken” by $dom(X)$ in r_4 . \square

To overcome unnecessary restrictions of strong safety in (20), *liberal domain-expansion safety* (lde-safety) has been introduced (15), which incorporates both syntactic and semantic properties of a program. The details of the notion are not necessary for this paper, except that all lde-safe programs have finite groundings with the same answer sets; we give here a brief overview.

Unlike strong safety, liberal de-safety is not a property of entire atoms but of *attributes*, i.e., pairs of predicates and argument positions. Intuitively, an attribute is lde-safe, if the number of different terms in an answer-set preserving grounding (i.e. a grounding which has the same answer sets if restricted to the positive atoms as the original program) is finite. A program is lde-safe, if all its attributes are lde-safe.

The notion of lde-safety is designed in an extensible fashion, i.e., such that several safety criteria can be easily integrated. For this we parametrize our definition of lde-safety by a *term bounding function* (TBF), which identifies variables in a rule that are ensured to have only finitely many instantiations in the answer set preserving grounding. Finiteness of the overall grounding follows then from the properties of TBFs.

For an ordinary predicate $p \in \mathcal{P}$, let $p \upharpoonright i$ be the i -th attribute of p for all $1 \leq i \leq ar(p)$. For an external predicate $\&g \in \mathcal{X}$ with input list \mathbf{X} in rule r , let $\&g[\mathbf{X}]_r \upharpoonright_T i$ with $T \in \{1, 0\}$ be the i -th input resp. output attribute of $\&g[\mathbf{X}]$ in r for all $1 \leq i \leq ar_T(\&g)$. For a ground program P , the range of an attribute is, intuitively, the set of ground terms which occur in the position of the attribute. Formally, for an attribute $p \upharpoonright i$ we have $range(p \upharpoonright i, P) = \{t_i \mid p(t_1, \dots, t_{ar(p)}) \in A(P)\}$; for an attribute $\&g[\mathbf{X}]_r \upharpoonright_T i$ we have $range(\&g[\mathbf{X}]_r \upharpoonright_T i, P) = \{x_i^T \mid \&g[\mathbf{x}^1](\mathbf{x}^0) \in EA(P)\}$, where $\mathbf{x}^s = x_1^s, \dots, x_{ar_s(\&g)}^s$.

We use the following monotone operator to compute by fixpoint iteration a finite subset of $grnd(P)$ for a program P :

$$G_P(P') = \bigcup_{r \in P} \{r\theta \mid \exists I \subseteq \mathcal{A}(P'), I \not\models \perp, I \models B^+(r\theta)\},$$

where $\mathcal{A}(P') = \{\mathbf{T}a, \mathbf{F}a \mid a \in A(P')\} \setminus \{\mathbf{F}a \mid a \leftarrow . \in P\}$ and $r\theta$ is the ground instance of r under variable substitution $\theta: \mathcal{V} \rightarrow \mathcal{C}$. Note that in this definition, I might be partial, but by convention we assume that all atoms which are not explicitly assigned to true are false. That is, G_P takes a ground program P' as input and returns all rules from $grnd(P)$ whose positive body is satisfied under some assignment over the atoms of Π' . Intuitively, the operator iteratively extends the grounding by new rules if they are possibly relevant for the evaluation, where relevance is in terms of satisfaction of the positive rule body under some assignment constructable over the atoms which are possibly derivable so far. Obviously, the least fixpoint $G_P^\infty(\emptyset)$ of this operator is a subset of $grnd(P)$; we will show that it is finite if P is lde-safe according to our new notion. Moreover, we will show that this grounding preserves all answer sets as all omitted rule instances have unsatisfied bodies anyway.

Example 32 Consider the following program P :

$$\begin{aligned} r_1: s(a). \quad r_2: dom(ax). \quad r_3: dom(axx). \\ r_4: s(Y) \leftarrow s(X), \&concat[X, x](Y), dom(Y). \end{aligned}$$

The least fixpoint of G_P is the following ground program:

$$\begin{aligned} r'_1: s(a). \quad r'_2: dom(ax). \quad r'_3: dom(axx). \\ r'_4: s(ax) \leftarrow s(a), \&concat[a, x](ax), dom(ax). \\ r'_5: s(axx) \leftarrow s(ax), \&concat[ax, x](axx), dom(axx). \end{aligned}$$

Rule r'_4 is added in the first iteration and rule r'_5 in the second.

Towards a definition of lde-safety, we say that a term in a rule is *bounded*, if the number of substitutions in $G_P^\infty(\emptyset)$ for this term is finite. This is abstractly formalized using *term bounding functions*.

Definition 21 (Term Bounding Function (TBF)) A term bounding function, denoted $b(P, r, S, B)$, maps a program P , a rule $r \in P$, a set S of (already safe) attributes, and a set B of (already bounded) terms in r to an enlarged set of (bounded) terms $b(P, r, S, B) \supseteq B$, such that every $t \in b(P, r, S, B)$ has finitely many substitutions in $G_P^\infty(\emptyset)$ if (i) the attributes S have a finite range in $G_P^\infty(\emptyset)$ and (ii) each term in $terms(r) \cap B$ has finitely many substitutions in $G_P^\infty(\emptyset)$.

Intuitively, a TBF receives a set of already bounded terms and a set of attributes that are already known to be lde-safe. Taking the program into account, the TBF then identifies and returns further terms which are also bounded.

The concept yields lde-safety of attributes and programs from the boundedness of variables according to a TBF. We provide a mutually inductive definition that takes the empty set of lde-safe attributes $S_0(P)$ as its basis. Then, each iteration step $n \geq 1$ defines first the set of bounded terms $B_n(r, P, b)$ for all rules r , and then an enlarged set of lde-safe attributes $S_n(P)$. The set of lde-safe attributes in step $n + 1$ thus depends on the TBF, which in turn depends on the domain-expansion safe attributes from step n .

Definition 22 (Liberal Domain-Expansion Safety) *Let b be a term bounding function. The set $B_n(r, P, b)$ of bounded terms in a rule $r \in P$ in step $n \geq 1$ is $B_n(r, P, b) = \bigcup_{j \geq 0} B_{n,j}(r, P, b)$ where $B_{n,0}(r, P, b) = \emptyset$ and for all $j \geq 0$, $B_{n,j+1}(r, P, b) = b(P, r, S_{n-1}(P), B_{n,j})$.*

The set of domain-expansion safe attributes $S_\infty(P) = \bigcup_{i \geq 0} S_i(P)$ of a program P is iteratively constructed with $S_0(P) = \emptyset$ and for $n \geq 0$:

- $p \upharpoonright i \in S_{n+1}(P)$ if for each $r \in P$ and atom $p(t_1, \dots, t_{ar(p)}) \in H(r)$, we have that term $t_i \in B_{n+1}(r, P, b)$, i.e., t_i is bounded;
- $\&g[\mathbf{X}]_r \upharpoonright_1 i \in S_{n+1}(P)$ if each \mathbf{X}_i is a bounded variable, or \mathbf{X}_i is a predicate input parameter p and $p \upharpoonright_1, \dots, p \upharpoonright_{ar(p)} \in S_n(P)$;
- $\&g[\mathbf{X}]_r \upharpoonright_0 i \in S_{n+1}(P)$ if and only if r contains an external atom $\&g[\mathbf{X}](\mathbf{Y})$ such that \mathbf{Y}_i is bounded, or $\&g[\mathbf{X}]_r \upharpoonright_1 1, \dots, \&g[\mathbf{X}]_r \upharpoonright_1 ar_1(\&g) \in S_n(P)$.

A program P is liberally domain-expansion (lde) safe, if it is safe and all its attributes are domain-expansion safe.

A detailed description of liberal safety is beyond the scope of this paper. However, it is crucial that each liberally domain-expansion safe HEX-program P is finitely restrictable, i.e., there is a finite subset P_g of $grnd_C(P)$ s.t. $\mathcal{AS}(P_g) = \mathcal{AS}(grnd_C(P))$. A concrete grounding algorithm GROUNDHEX is given in (15); we use GROUNDHEX(P) in this article to refer to a finite grounding of P that has the same answer sets.

References

- [1] Eyal Amir and Sheila A. McIlraith. Partition-based logical reasoning for first-order and propositional theories. *Artificial Intelligence*, 162(1-2):49–88, 2005.
- [2] Seif El-Din Bairakdar, Minh Dao-Tran, Thomas Eiter, Michael Fink, and Thomas Krennwallner. The DMCS solver for distributed nonmonotonic multi-context systems. In *European Conference on Logics in Artificial Intelligence (JELIA)*, pages 352–355. Springer, 2010.
- [3] Selen Basol, Ozan Erdem, Michael Fink, and Giovambattista Ianni. HEX programs with action atoms. In *Technical Communications of the International Conference on Logic Programming (ICLP)*, pages 24–33, 2010.
- [4] Markus Bögl, Thomas Eiter, Michael Fink, and Peter Schüller. The MCS-IE system for explaining inconsistency in multi-context systems. In *European Conference on Logics in Artificial Intelligence (JELIA)*, pages 356–359, 2010.

- [5] Gerd Brewka and Thomas Eiter. Equilibria in heterogeneous nonmonotonic multi-context systems. In *AAAI Conference on Artificial Intelligence*, pages 385–390. AAAI Press, 2007.
- [6] Francesco Calimeri, Susanna Cozza, and Giovambattista Ianni. External sources of knowledge and value invention in logic programming. *Annals of Mathematics and Artificial Intelligence*, 50(3–4):333–361, 2007.
- [7] Francesco Calimeri, Susanna Cozza, Giovambattista Ianni, and Nicola Leone. Computable functions in ASP: Theory and implementation. In *ICLP, LNCS*, pages 407–424. Springer, 2008.
- [8] Francesco Calimeri, Michael Fink, Stefano Germano, Giovambattista Ianni, Christoph Redl, and Anton Wimmer. AngryHEX: an artificial player for angry birds based on declarative knowledge bases. In *National Workshop and Prize on Popularize Artificial Intelligence*, pages 29–35, 2013.
- [9] Edmund Clarke, Orna Grumberg, Somesh Jha, Yuan Lu, and Helmut Veith. Counterexample-guided abstraction refinement for symbolic model checking. *J. ACM*, 50(5):752–794, 2003.
- [10] Minh Dao-Tran, Thomas Eiter, and Thomas Krennwallner. Realizing default logic over description logic knowledge bases. In *Symbolic and Quantitative Approaches to Reasoning with Uncertainty*, pages 602–613. Springer, 2009.
- [11] Phan Minh 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.
- [12] Thomas Eiter, Michael Fink, Giovambattista Ianni, Thomas Krennwallner, and Peter Schüller. Pushing efficient evaluation of HEX programs by modular decomposition. In *International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR)*, pages 93–106, 2011.
- [13] Thomas Eiter, Michael Fink, and Thomas Krennwallner. Decomposition of Declarative Knowledge Bases with External Functions. In *International Joint Conference on Artificial Intelligence (IJCAI)*, pages 752–758. AAAI Press, 2009.
- [14] Thomas Eiter, Michael Fink, Thomas Krennwallner, and Christoph Redl. Conflict-driven ASP solving with external sources. *Theory and Practice of Logic Programming*, 12(4-5):659–679, 2012.
- [15] Thomas Eiter, Michael Fink, Thomas Krennwallner, and Christoph Redl. Domain expansion for ASP-programs with external sources. Technical Report INFSYS RR-1843-14-02, Institut für Informationssysteme, Technische Universität Wien, A-1040 Vienna, Austria, 2014.
- [16] Thomas Eiter, Michael Fink, Thomas Krennwallner, Christoph Redl, and Peter Schüller. Efficient HEX-program evaluation based on unfounded sets. *Journal of Artificial Intelligence Research*, 49:269–321, 2014.
- [17] Thomas Eiter, Georg Gottlob, and Heikki Mannila. Disjunctive datalog. *ACM Transactions on Database Systems*, 22(3):364–418, 1997.
- [18] Thomas Eiter, Giovambattista Ianni, Thomas Lukasiewicz, Roman Schindlauer, and Hans Tompits. Combining answer set programming with description logics for the semantic web. *Artificial Intelligence*, 172(12-13):1495–1539, 2008.

- [19] Thomas Eiter, Giovambattista Ianni, Roman Schindlauer, and Hans Tompits. A Uniform Integration of Higher-Order Reasoning and External Evaluations in Answer-Set Programming. In *International Joint Conference on Artificial Intelligence (IJCAI)*, pages 90–96. Professional Book Center, 2005.
- [20] Thomas Eiter, Giovambattista Ianni, Roman Schindlauer, and Hans Tompits. Effective integration of declarative rules with external evaluations for semantic-web reasoning. In *European Semantic Web Conference (ESWC)*, pages 273–287. Springer, 2006.
- [21] Wolfgang Faber, Nicola Leone, and Gerald Pfeifer. Recursive aggregates in disjunctive logic programs: Semantics and complexity. In *European Conference on Logics in Artificial Intelligence (JELIA)*, pages 200–212. Springer, 2004.
- [22] M. Gebser, M. Ostrowski, and T. Schaub. Constraint answer set solving. In *International Conference on Logic Programming (ICLP)*, pages 235–249. Springer, 2009.
- [23] Martin Gebser, Roland Kaminski, Benjamin Kaufmann, and Torsten Schaub. Clingo = ASP + control: Preliminary report. *CoRR*, abs/1405.3694, 2014.
- [24] Martin Gebser, Benjamin Kaufmann, and Torsten Schaub. Solution enumeration for projected boolean search problems. In *Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems (CPAIOR)*, pages 71–86. Springer, 2009.
- [25] Martin Gebser, Benjamin Kaufmann, and Torsten Schaub. Conflict-driven answer set solving: From theory to practice. *Artificial Intelligence*, 187–188:52–89, 2012.
- [26] M. Gelfond and V. Lifschitz. The Stable Model Semantics for Logic Programming. In R. Kowalski and K. Bowen, editors, *Logic Programming: Proceedings of the 5th International Conference and Symposium*, pages 1070–1080. MIT Press, 1988.
- [27] Michael Gelfond and Vladimir Lifschitz. Classical negation in logic programs and disjunctive databases. *New Generation Computing*, 9(3/4):365–386, 1991.
- [28] Giray Havur, Guchan Ozbilgin, Esra Erdem, and Volkan Patoglu. Geometric Rearrangement of Multiple Movable Objects on Cluttered Surfaces: A Hybrid Reasoning Approach. In *International Conference on Robotics and Automation (ICRA)*, pages 445–452, 2014.
- [29] Robert Hoehndorf, Frank Loebe, Janet Kelso, and Heinrich Herre. Representing default knowledge in biomedical ontologies: Application to the integration of anatomy and phenotype ontologies. *BMC Bioinformatics*, 8(1):377, 2007.
- [30] Tomi Janhunnen, Emilia Oikarinen, Hans Tompits, and Stefan Woltran. Modularity Aspects of Disjunctive Stable Models. *Journal of Artificial Intelligence Research*, 35:813–857, 2009.
- [31] Matti Järvisalo, Emilia Oikarinen, Tomi Janhunnen, and Ilkka Niemelä. A module-based framework for multi-language constraint modeling. In *Logic Programming and Nonmonotonic Reasoning (LPNMR)*, pages 155–168, 2009.
- [32] O. Lassila and R.R. Swick. Resource description framework (RDF) model and syntax specification, 1999. <http://www.w3.org/TR/1999/REC-rdf-syntax-19990222>.

- [33] Yuliya Lierler and Mirosław Truszczyński. Modular answer set solving. In *Late-Breaking Developments in the Field of Artificial Intelligence, Bellevue, Washington, USA, July 14-18, 2013*, volume WS-13-17 of *AAAI Workshops*. AAAI, 2013.
- [34] V. Lifschitz and H. Turner. Splitting a Logic Program. In *Proceedings ICLP-94*, pages 23–38, Santa Margherita Ligure, Italy, 1994. MIT-Press.
- [35] Thomas Linke. Graph Theoretical Characterization and Computation of Answer Sets. In *International Joint Conference on Artificial Intelligence (IJCAI)*, pages 641–645, 2001.
- [36] Thomas Linke and Vladimir Sarsakov. Suitable graphs for answer set programming. In Franz Baader and Andrei Voronkov, editors, *LPAR*, volume 3452 of *Lecture Notes in Computer Science*, pages 154–168. Springer, 2004.
- [37] Alessandro Mosca and Diego Bernini. Ontology-driven geographic information system and dlvh reasoning for material culture analysis. In *Italian Workshop RiCeRcA at ICLP*, 2008.
- [38] Ilkka Niemelä. Logic programming with stable model semantics as constraint programming paradigm. *Annals of Mathematics and Artificial Intelligence*, 25(3–4):241–273, 1999.
- [39] Emilia Oikarinen and Tomi Janhunen. Achieving compositionality of the stable model semantics for smodels programs. *TPLP*, 8(5-6):717–761, 2008.
- [40] Max Ostrowski and Torsten Schaub. ASP modulo CSP: the clingcon system. *Theory and Practice of Logic Programming (TPLP)*, 12(4-5):485–503, 2012.
- [41] Alessandro Dal Palù, Agostino Dovier, Enrico Pontelli, and Gianfranco Rossi. Gasp: Answer set programming with lazy grounding. *Fundamenta Informaticae*, 96(3):297–322, 2009.
- [42] Simona Perri, Francesco Ricca, and Marco Sirianni. A parallel ASP instantiator based on DLV. In *Declarative Aspects of Multicore Programming (DAMP'10)*, LNCS, pages 73–82. Springer, 2010.
- [43] Axel Polleres. From SPARQL to rules (and back). In *International Conference on World Wide Web (WWW)*, pages 787–796. ACM, 2007.
- [44] Teodor C. Przymusiński. Stable semantics for disjunctive programs. *New Generation Computing*, 9:401–424, 1991.
- [45] Teodor C. Przymusiński. On the Declarative Semantics of Deductive Databases and Logic Programs. In J. Minker, editor, *Foundations of Deductive Databases and Logic Programming*, pages 193–216. Morgan Kaufman, 1988.
- [46] Christoph Redl. *Answer Set Programming with External Sources: Algorithms and Efficient Evaluation*. PhD thesis, Vienna University of Technology, 2014.
- [47] K.A. Ross. Modular Stratification and Magic Sets for Datalog Programs with Negation. *Journal of the ACM*, 41(6):1216–1267, 1994.
- [48] Roman Schindlauer. *Answer Set Programming for the Semantic Web*. PhD thesis, Vienna University of Technology, Vienna, Austria, 2006.

- [49] Peter Schüller. *Inconsistency in Multi-Context Systems: Analysis and Efficient Evaluation*. PhD thesis, Vienna University of Technology, Vienna, Austria, 2012.
- [50] Peter Schüller, Volkan Patoglu, and Esra Erdem. A Systematic Analysis of Levels of Integration between Low-Level Reasoning and Task Planning. In *Workshop on Combining Task and Motion Planning at IEEE International Conference on Robotics and Automation (ICRA)*, 2013.
- [51] Yi-Dong Shen, Kewen Wang, Thomas Eiter, Michael Fink, Christoph Redl, Thomas Krennwallner, and Jun Deng. FLP answer set semantics without circular justifications for general logic programs. *Artificial Intelligence*, 213:1–41, 2014.
- [52] Shahab Tasharrofi and Eugenia Ternovska. A semantic account for modularity in multi-language modelling of search problems. In *International Symposium on Frontiers of Combining Systems (FroCoS)*, pages 259–274, 2011.
- [53] Yisong Wang, Jia-Huai You, Li-Yan Yuan, Yi-Dong Shen, and Mingyi Zhang. The loop formula based semantics of description logic programs. *Theor. Comput. Sci.*, 415:60–85, 2012.
- [54] Jesia Zakraoui and Wolfgang L. Zagler. A method for generating CSS to improve web accessibility for old users. In *Int. Conf. on Computers Helping People with Special Needs (ICCHP)*, pages 329–336, 2012.
- [55] Hande Zirtiloğlu and Pinar Yolum. Ranking semantic information for e-government: complaints management. In *International Workshop on Ontology-supported business intelligence (OBI)*. ACM, 2008.