



The Hexlite Solver

Lightweight and Efficient Evaluation of HEX Programs

Peter Schüller^(✉) 

Institut für Logic and Computation, Knowledge-Based Systems Group,
Technische Universität Wien, Vienna, Austria
`ps@kr.tuwien.ac.at`

Abstract. HEXLITE is a lightweight solver for the HEX formalism which integrates Answer Set Programming (ASP) with external computations. The main goal of HEXLITE is efficiency and simplicity, both in implementation as well as in installation of the system. We define the Pragmatic HEX Fragment which permits to partition external computations into two kinds: those that can be evaluated during the program instantiation phase, and those that need to be evaluated during the answer set search phase. HEXLITE is written in PYTHON and suitable for evaluating this fragment with external computations that are realized in PYTHON. Most performance-critical tasks are delegated to the PYTHON module of CLINGO. We demonstrate that the Pragmatic HEX Fragment is sufficient for many use cases and that it permits HEXLITE to have superior performance compared to the DLVHEX system in relevant application scenarios.

1 Introduction

The HEX formalism [8] facilitates the combination of logic programming and external computations in other programming paradigms, and facilitates the integration of logical reasoning with diverse other reasoning methods such as motion planning [20], description logics [14], or sub-symbolic reasoning [17].

Differently from externals in GRINGO and the PYTHON interface of CLINGO, the HEX formalism provides uniform and generic syntax and semantics for external computations that influence (a) the instantiation of the program (by performing value invention), and (b) the solving process (by computing truth values relative to interpretations).

Computing HEX semantics requires the evaluation of external computations both during grounding and during search, in an interleaved fashion. The main solver implementation for the HEX formalism is the DLVHEX system [13]. While DLVHEX implements the full HEX language, it performs a lot of analysis and preprocessing to be able to deal with all eventualities of combinations of external computations, which makes it unnecessarily slow in several relevant application scenarios.

In order to obtain performance when evaluating HEX programs, we here present a new HEX solver that is lightweight and efficient, at the cost of handling only a fragment of HEX, which is sufficient for many applications.

In this work, we make the following contributions.

- We define the Pragmatic HEX Fragment (PHF) that permits to separate external computations into those that can be evaluated completely during instantiation and those that can be evaluated completely during search. We show several application scenarios where the PHF is sufficient.
- We describe and provide the HEXLITE solver that provides lightweight and efficient evaluation machinery for the PHF. The solver is implemented in PYTHON and uses the CLINGO PYTHON API as a backend for ASP grounding and search. HEXLITE rewrites both classes of external atoms in different ways before passing the rules to CLINGO for evaluation. As a main benefit of this architecture, HEXLITE supports the full ASP input language of CLINGO without the need for dedicated code that supports weak constraints, choice rules, aggregates, expansion terms, and builtin arithmetics.
- We experimentally compare the HEXLITE solver with DLVHEX on two application scenarios that gave rise to the development of HEXLITE: cost-based abduction [30] and RDF processing [18]. Our experiments show, that HEXLITE performs better than DLVHEX in these applications. As HEXLITE uses the same PYTHON API as DLVHEX, we use the same plugin with both solvers, which makes the comparison very realistic.

HEXLITE can be installed via CONDA or PIP and is available as open source.¹

2 Preliminaries

We give syntax and semantics of the HEX formalism [8, 16] which generalizes logic programs under answer set semantics [24] with external computations.

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. Usually, elements from \mathcal{X} and \mathcal{C} are denoted with first letter in upper case and lower case, respectively; while elements from \mathcal{G} are prefixed with ‘&’. Elements from $\mathcal{C} \cup \mathcal{X}$ are called *terms*. An (ordinary) *atom* is a tuple $p(Y_1, \dots, Y_n)$ where $p \in \mathcal{C}$ is a predicate name and Y_1, \dots, Y_n are terms and $n \geq 0$ is the *arity* of the atom. An atom is *ground* if all its terms are constants. An *external atom* is of the form $\&g[Y_1, \dots, Y_n](X_1, \dots, X_m)$, 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 input and output lists have fixed lengths $in(\&g) = n$ and $out(\&g) = m$. With each term Y_i in the input list, $1 \leq i \leq n$, we associate a type $t_i \in \{\mathbf{cons}\} \cup \mathbb{N}$. We call the term *constant input* iff $t_i = \mathbf{cons}$, otherwise we call it *predicate input of arity* t_i .

¹ <https://github.com/hexhex/hexlite>.

A rule r is of the form $\alpha_1 \vee \dots \vee \alpha_k \leftarrow \beta_1, \dots, \beta_n, \text{not } \beta_{n+1}, \dots, \text{not } \beta_m$ with $m, k \geq 0$ 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 $B^-(r) = \{\beta_{n+1}, \dots, \beta_m\}$. A rule r is a *constraint* if $H(r) = \emptyset$; 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.

A *HEX program* is a finite set P of rules. We call a program P *ordinary* (resp., *nondisjunctive*) if all its rules are ordinary (resp., nondisjunctive). Note that we here assume that programs have no *higher-order* atoms (i.e., atoms of the form $Y_0(Y_1, \dots, Y_n)$ where $Y_0 \in \mathcal{X}$) because HEX-programs with higher-order atoms can easily be rewritten to HEX-programs without higher-order atoms [8].

A comprehensive introduction to HEX is given in [18].

2.2 Semantics

Given a rule r , the grounding $grnd(r)$ of r is obtained by systematically replacing all variables with constants from \mathcal{C} . Given a HEX-program P , the *Herbrand base* HB_P of 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 $grnd(P)$ of P is given by $grnd(P) = \bigcup_{r \in P} grnd(r)$. Importantly, the set of constants \mathcal{C} that is 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 the semantics of the program.

Extensional Semantics [8,16] of external atoms are defined as follows: we associate a $(n + 1)$ -ary extensional evaluation function $F_{\&g}$ with every external predicate name $\&g \in \mathcal{G}$. Given an interpretation $I \subseteq HB_P$ and a ground input tuple (x_1, \dots, x_m) , $F_{\&g}(I, y_1, \dots, y_n)$ returns a set of ground output tuples (x_1, \dots, x_m) . The external computation is *restricted* to depend (a) for constant inputs, i.e., $t_i = \mathbf{cons}$, only on the constant value of y_i ; and (b) for predicate inputs, i.e., $t_i \in \mathbb{N}$, only on the extension of predicate y_i of arity t_i in I .²

An interpretation $I \subseteq HB_P$ is a *model* of an atom a , denoted $I \models a$ if a is an ordinary atom and $a \in I$. I is a model of a ground external atom $a = \&g[y_1, \dots, y_n](x_1, \dots, x_m)$ if $(x_1, \dots, x_m) \in F_{\&g}(I, y_1, \dots, y_n)$. Given a ground rule r , $I \models H(r)$ if $I \models a$ for some $a \in H(r)$; $I \models B(r)$ if $I \models a$ for all $a \in B^+(r)$ and $I \not\models a$ for all $a \in B^-(r)$; and $I \models r$ if $I \models H(r)$ whenever $I \models B(r)$. Given a HEX-program P , $I \models P$ if $I \models r$ for all $r \in grnd(P)$; the *FLP-reduct* [21] of P with respect to $I \subseteq HB_P$, denoted fP^I , is the set of all $r \in grnd(P)$ such that $I \models B(r)$; $I \subseteq HB_P$ is an *answer set* of P if I is a minimal model of fP^I , and we denote by $\mathcal{AS}(P)$ the set of all answer sets of P .

3 The Pragmatic HEX Fragment (PHF)

We next define a fragment of HEX that permits to separate external computations into two classes: grounding-relevant and solving-relevant.

² Formally, this is the set $\{y_i(v_1, \dots, v_{t_i}) \in I\}$.

Definition 1. A HEX-program P is in the Pragmatic HEX Fragment (PHF) iff each external atom of the form $\&g[Y_1, \dots, Y_n](X_1, \dots, X_m)$ with type signature t_1, \dots, t_n satisfies one of the following conditions:

- (G) $t_i = \mathbf{cons}$ for all i , $1 \leq i \leq n$; or
- (S) $m = 0$ and there is at least one type t_i , $1 \leq i \leq n$, such that $t_i \in \mathbb{N}$.

Type (G). External atoms that satisfy condition (G) can be *evaluated during instantiation* of program P because their computation does not depend on I . These external computations can perform value invention: they can produce constants in the output tuple X_1, \dots, X_m that do not exist in P . In particular, output terms can recursively define input terms of the same external atom.

Example 1. We can use an external atom of the form $\&rdf[U](S, P, O)$ of type (G) to access a RDF [26] triple stores [18]. The function $F_{\&rdf}(I, U)$ returns all tuples (S, P, O) that are obtained from the RDF graph accessible at URI U . Intuitively, this external computation imports the RDF graph into the HEX program and provides all its constants to the instantiation process. The computation does not depend on I and returns arbitrary strings (value invention).

Type (S). External atoms that satisfy condition (S) have an empty output tuple and can therefore not produce any output apart from their own truth value. This makes it possible to instantiate rules that contain such external atoms, without performing the associated external computation; the external computation needs to be performed only *during the answer set search* phase, when the interpretation I is available.

Example 2. We can use external atoms of the form $\&transitive[p]()$ of type (S), with $t_1 = 2$ the arity of predicate p , to verify whether $p/2$ is transitive in the interpretation [30]. The function $F_{\&transitive}(I, p)$ returns the empty tuple if the extension of p in I is a transitive relation. Otherwise, it returns no tuple.

External atoms of type (S) have the possibility to create nogoods that relate the truth value of the ground external atom with parts of I that are relevant for computing that truth value. This feature, which also exists in DLVHEX, can be used for increasing evaluation performance by guiding the solver towards answer set candidates that are compatible with external computation results.

3.1 Properties

Clearly, the two classes (G) and (S) of external atoms are mutually exclusive. Those external atoms of type (G) that have no output terms (i.e., $m = 0$) could be evaluated during the solving process because their evaluation is not necessary for instantiating the ground program. However, we found it useful to evaluate (and therefore eliminate) as many external atoms as possible already during program instantiation.

Moreover, there are external atoms that fall neither into class (G) nor into class (S), for example an external atom $\&sum[pred](X)$ that has one predicate

input $pred$ of arity 1 and realizes a summation aggregate with extensional evaluation function $F_{\&sum}(I, pred) := \{(X)\}$ where $X = \Sigma\{x \mid p(x) \in I\}$.

External atoms of type (G) can be evaluated (and eliminated) during instantiation of the program as shown in the following proposition.

Proposition 1. *Given a ground HEX-program P in PHF, an equivalent program P' can be produced by (1) omitting rules that contain an external atom a of type (G) where (i) a is in a positive literal and $\emptyset \not\models a$; or (ii) a is in a negative literal and $\emptyset \models a$; and (2) omitting all other external atoms of type (G).*

Proof (sketch). A ground external atom a of type (G) has only constant input, therefore $I \models a$ is independent from the value of I . Hence, rules where (1) applies can never obtain a satisfied body due to a while rules where (2) applies can never obtain a non-satisfied body due to a .

External atoms of type (S) can be handled the same way as in DLVHEX [8, 13], therefore we here do not provide formal results about them.³

External atoms outside the PHF fragment can be processed with Liberal Safety [10]: it permits automatic verification of finite instantiation of a HEX program in the presence of cyclic dependencies among external atoms that perform value invention, depend on the answer set candidate, and have certain (semantic) properties. Opposed to the methodology of Liberal Safety, HEXLITE delegates finiteness of instantiation to the programmer (as also done, e.g., by GRINGO).

3.2 Amenable Application Scenarios

External computations of type (S) can pass nogoods to the solver that describe how their truth value depends on the interpretation.

Constraint Answer Set Programming (CASP) [27] has been realized in HEX [29] using one external atom $\&check$ of type (S). Application scenarios in [29] use external atoms of type (G) for SQL querying: $\&sql[Query](AnswerTuple)$, where $Query$ can be a fixed string or defined using external atoms. High-level *planning for robotics* has been interleaved with low-level motion planning using HEX [20], where external atoms for motion planning are either of type (S) or of type (G) and there is no value invention. HEX-programs with *existential quantification* (HEX[∃]) [9] as well as HEX-programs with function symbols use only external atoms of type (G) to perform tasks related to Skolemization, similar to what is presented in Sect. 5.1. The MCS-IE system for *explaining inconsistency in Multi-Context Systems* [4] is implemented in HEX and uses only external atoms of type (S). Two further application scenarios are *abduction* and *RDF processing*, shown in detail in Sects. 5.1 and 5.2.

³ DLVHEX replaces them with an ordinary replacement atom, guesses truth of replacement atoms with extra rules, and accepts only answer set candidates I where guessed truth values correspond with external computations wrt. I , see Sect. 4.

4 Hexlite Solver Design and Architecture

Principles. The design of HEXLITE followed several guiding principles.

- Delegation: delegate as much as possible to the backend solver (currently CLINGO).
- Separation: deal with external atoms either during grounding or during solving in order to avoid multiple grounding passes.
- Programmer Responsibility: delegate the responsibility for finite instantiation to the programmer (as in GRINGO).

Delegation reduces computational overhead and duplication of code that already exists in the ASP backend. As the main consequence of this principle, HEXLITE performs no safety check and no syntax check, not even thorough parsing of the input. Instead, a shallow representation of the input program is created. This representation is sufficient for rewriting rules and external atoms for subsequent evaluation. As a consequence of Delegation, unsafe variables are detected only by GRINGO because safety checking is not required for HEXLITE rewriting. A second consequence of Delegation is, that HEXLITE has no internal representation of the current answer set candidate; instead, the CLASP PYTHON API is used to directly access the (partial) model within CLASP. Moreover, optimization is handled transparently within CLASP.

Separation is the basis for defining the PHF and contributes to the small implementation of HEXLITE, because evaluation follows the common structure of ASP solvers with external atoms that are relevant either for grounding or for solving. (Currently, HEXLITE is implemented using 3,300 lines of PYTHON code, which includes the shallow parser, FLP checker, and code comments.) A HEX program can be split into Evaluation Units, i.e., non-ground program modules that never mutually depend on one another [8]. While DLVHEX interleaves grounding and search over multiple Evaluation Units, HEXLITE always uses a single Evaluation Unit.

Programmer Responsibility is a principle from GRINGO: the burden of ensuring a finite instantiation is put on the programmer and not verified by preprocessing. This is the opposite of the philosophy followed in the DLVHEX solver where various safety notions such as Domain Expansion Safety and Liberal Domain Expansion Safety [11] are defined and also checked by the solver, depending on properties of external computations. The upside of not checking this in the solver is decreased preprocessing effort, the downside is that HEXLITE (just like GRINGO) will not complain about the program ‘ $p(0) \leftarrow . p(s(X)) \leftarrow p(X).$ ’ but starts to instantiate it and will exhaust the available memory due to its infinite instantiation. In HEX programs, external computations can cause infinite instantiation, but the programmer of external computations might take specific measures to prevent infinite instantiation (as shown in Sect. 5.1). Therefore, delegating responsibility to the user instead of performing costly verifications can be an advantage.

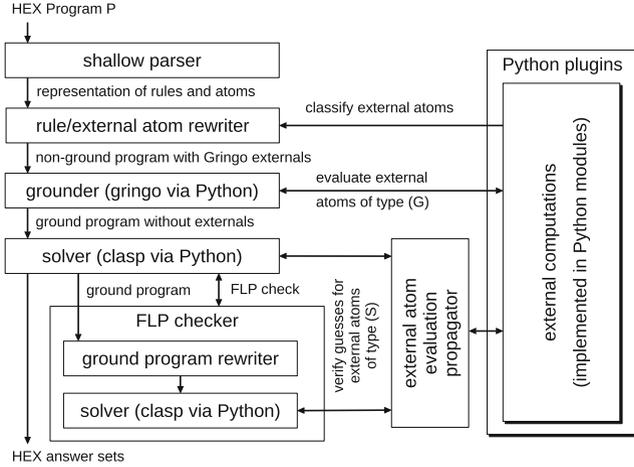


Fig. 1. Architecture of the HEXLITE system and its interaction with the CLINGO Python library and plugins for external computations.

Architecture. Figure 1 shows the architecture of HEXLITE. The input program P is analyzed with a *shallow parser*, followed by the *rewriter* module that rewrites only those parts of rules that contain external atoms. The *rewriter* accesses external computations for obtaining their type and creates the following GRINGO-compatible rules for each rule $r \in P$.

- (i) a rule r' where each external atom a of the form $\&g[Y_1, \dots, Y_n](X_1, \dots, X_m)$
 - (a) is replaced by a GRINGO external of the form $(X_1, \dots, X_m) = @g(Y_1, \dots, Y_n)$ if a is of type (G);
 - (b) is replaced by a *replacement atom* of the form $e_{\&g}(Y_1, \dots, Y_n)$ if a is of type (S); and⁴
- (ii) for each replacement atom created in (b) above a rule r'' of the form

$$e_{\&g}(Y_1, \dots, Y_n) \vee \bar{e}_{\&g}(Y_1, \dots, Y_n) \leftarrow B'$$

where $B' = \{\beta \in B(r) \mid \beta \text{ shares variables with } Y_1, \dots, Y_n\}$.

(DLVHEX uses a similar rewriting, see [15, 19].) After rewriting, GRINGO is used for instantiating the rewritten program with a PYTHON context that evaluates external computations of the form $(X_1, \dots, X_m) = @g(Y_1, \dots, Y_n)$ by assigning the tuple (X_1, \dots, X_m) all values returned by the external computation. The resulting ordinary ground program is passed to a CLASP instance that has been prepared with (i) a custom PYTHON propagator,⁵ (ii) a ground program observer that collects the ground program for usage in the FLP checker, and (iii) an

⁴ In that case $m = 0$ so the replacement atom does not include X_1, \dots, X_m .

⁵ A propagator is a program module that interfaces with the search process of CLASP and can (a) infer truth values and (b) add ground clauses based on a partial model.

`on_model` callback (not shown in the figure) which calls the FLP checker and, after a successful check, outputs the answer set without auxiliary elements.

The custom PYTHON propagator compares truth values of replacement atoms with the results of their corresponding external computations. For failed checks, nogoods are created to prevent future failed checks. Moreover, user-defined nogoods can be provided by plugins to further guide the search (see also the description of external atoms in Sects. 5.1 and 5.2). The FLP checker is necessary for HEX programs with positive loops over external atoms and prevents answer sets with self-founded truth values, see also [21]. It also ensures that non-monotonic aggregates are evaluated correctly (CLINGO does not support FLP semantics).⁶

HEXLITE uses the same PYTHON API as DLVHEX, therefore migrating plugins from one solver to the other is easy. None of the benchmarks we used in evaluations has positive loops over external atoms of type (S), therefore we deactivated the FLP checker for all experiments (both for DLVHEX and for HEXLITE).

5 Experimental Evaluation

For experimental evaluation of HEXLITE, we tried it out with two application domains: variants of cost-based abduction [30] and the RDF plugin [18].

We choose these domains because they both contain both types of external atoms, and because they were part of the inspiration for developing HEXLITE.

5.1 Cost-Based Abduction Benchmark

Cost-based abduction consists of 50 instances over the ACCEL natural language story understanding benchmark [28]. The full benchmark is available online.^{7,8} ACCEL uses two external predicates: *&invent* for flexible Skolemization, and *&transitive* as described in Example 2 for ensuring transitivity of a guessed relation over many elements.

External predicate *&invent* is used for ensuring finite instantiation in the presence of existential variables in rule heads.⁹ Given a rule R with an existential variable B in the head and universal variables A_1, \dots, A_k in the body, instead of replacing B with Skolem term $s(A_1, \dots, A_k)$ we add *&invent* $[R, c_B, A_1, \dots, A_k](B)$ to the rule body and define $F_{\&invent}$ such that a finite set of values for B is

⁶ The FLP check implemented in HEXLITE is described in Proposition 1 in [8]. The FLP check can be deactivated if it is not required (this is another example of Programmer Responsibility). The custom PYTHON propagator is re-used in the FLP checker.

⁷ <https://bitbucket.org/knowlp/asp-fo-abduction>.

⁸ To permit a fairer comparison, we used only objective functions CARD and COH (DLVHEX is incompatible with objective function WA) and we removed all facts of the form `comment(...)`, which served only an informational purpose (DLVHEX is significantly slower if these facts are included).

⁹ This generalizes the termination mechanism for reasoning as it was implemented in the original ACCEL reasoner [28].

invented, independent from cycles in the program. For our evaluation, we used two variants of this benchmark:

SK/P¹ variant. Value invention is blocked if at least one parent term is invented: $F_{\&invent}^{P^1}(I, R, V, A_1, \dots, A_k)$ returns the tuple $(s(R, V, A_1, \dots, A_k))$ if no A_i , $1 \leq i \leq k$ is of the form $s(\dots)$, otherwise it returns no tuple.

SK/G¹ variant. Value invention is blocked if at least one parent term was invented by the same ground external predicate: $F_{\&invent}^{G^1}(I, R, V, A_1, \dots, A_k)$ returns the tuple $(s(R, V, A_1, \dots, A_k))$ if no A_i , $1 \leq i \leq k$ has a sub-term of the form $s(R, V, \dots)$, otherwise it returns no tuple.

Intuitively, SK/P¹ invents only values that have non-invented *parents*, and SK/G¹ invents only a single *generation* of values in each rule. This method of ensuring finite instantiation is orthogonal to guarding domains predicates that ensure a finite chase in Datalog with existential quantifiers [5]: value invention is forced to be finite independent from the structure of the program. For details and examples of the finite instantiation property, we refer to Sect. 4.1 in [30].

External predicate $\&transitive$ appears in a single constraint of the form $\leftarrow \text{not } \&transitive[eq]()$. Each time this external computation evaluates to false, it creates nogoods that provide the reason for intransitivity (i.e., a triple of literals $eq(A, B)$, $eq(B, C)$, $\neg eq(A, C)$) to the solver to assist the search for a transitive relation. For more details see Sect. 4.2 in [30].

Importantly, $F_{\&invent}$ does not use I and therefore external atoms of the form $\&invent\dots$ are of type (G). Moreover, external predicate $\&transitive$ produces no values in output tuples (i.e., $m = 0$) and it has type $t_1 = 2$ (it uses the extension of a binary predicate p) therefore it is of type (S).

5.2 RDF Benchmark

The RDF plugin realizes the RDF triple external atom as described in Example 1. We here extend the original application [18] with a second external atom of type (S). We experiment with the COLINDA [31] knowledge graph which contains 150,000 triples about conferences.¹⁰ We perform the following three reasoning problems.

Import. We simply import all triples using the program

```
explore(".../colinda.rdf") ← .
triple_at(S, P, O) ← &rdf[What](S, P, O), explore(What).
```

which yields a single answer set with around 150,000 atoms.

Vegas. We are interested in names of all conferences in Las Vegas. The program

```
explore(".../colinda.rdf") ← . location("http://sws.geonames.org/3635260/") ← .
conference(Conf) ← explore(G), location(Loc),
&rdf[G](Conf, "http://swrc.ontoware.org/ontology#location", Loc).
```

¹⁰ Dataset retrieved from <https://old.datahub.io/dataset/colinda>.

$$\begin{aligned} title(Title) &\leftarrow explore(G), conference(Conf), \\ &\&rdf[G](Conf, "http://swrc.ontoware.org/ontology\#eventTitle", Title). \end{aligned}$$

yields a single answer set with 330 atoms, 164 of them containing the titles conferences in COLINDA that took place in Las Vegas (encoded as ID 3635260).

Marathon. This is an optimization problem where we are interested in making a conference marathon for visiting the maximum possible number of cities in France over two weeks. This is encoded in the following HEX program.

$$\begin{aligned} explore(".../colinda.rdf") &\leftarrow . \quad country("France") \leftarrow . \\ location(Loc) &\leftarrow explore(G), country(C), \\ &\&rdf[G](Loc, "http://www.geonames.org/ontology\#countryName", C). \\ conference(Conf, Loc) &\leftarrow explore(G), location(Loc), \\ &\&rdf[G](Conf, "http://swrc.ontoware.org/ontology\#location", Loc). \\ in(Conf) \vee out(Conf) &\leftarrow conference(Conf, Loc). \\ covered(Loc) &\leftarrow conference(Conf, Loc), in(Conf). \\ \leftarrow location(Loc), not \textit{covered}(Loc). \quad [1, Loc] \\ date(Date) &\leftarrow explore(G), in(Conf), \\ &\&rdf[G](Conf, "http://swrc.ontoware.org/ontology\#startDate", Date). \\ \leftarrow not \&dates_span_days[date, 14]. \end{aligned}$$

This encoding represents locations in France in *location*, conferences and their locations (if in France) in *conference*, and performs a guess (*in*) selecting relevant conferences. Those conferences that are selected define which location is covered. We maximize coverage by incurring a cost of 1 for each location that is not covered by means of the weak constraint (\leftarrow). Furthermore, we extract dates of covered conferences in *date* and use an external atom to indicate whether the dates lie within a 14 day period.

Function $F_{\&dates_span_days}(I, p, d)$ is defined to return an empty tuple iff all dates X with $p(X) \in I$ are within d days of one another. The external computation of *dates_span_days* guides the search by providing nogoods for all pairs of dates (X, X') with $p(X), p(X') \in I$ and more than d days between X and X' .

Evaluating this program yields an answer set with 29 locations and 404 conferences in France, and an optimal selection of 6 conferences (in 6 distinct cities) which start between the 12th and the 23rd of March 2012.

5.3 Experimental Setup

We performed experiments on a computer with an Intel(R) i5-3450 CPU with 4 cores and 16 GB RAM running Linux. For the Abduction Benchmark we limited memory consumption to 5 GB and execution time to 300 s (5 min). For the RDF Benchmark we limited execution time to 1800 s (30 min). We never executed more than 2 runs in parallel and we used non-parallel computation mode

Table 1. RDF benchmark: results of evaluating DLVHEX and HEXLITE on the COLINDA conference knowledge graph.

Problem	Engine	Time (s)	Space (MB)	Result	External computations		
					Calls (#)	Time (s)	Learned (#)
Import	DLVHEX	1,800	4,321	TO	1	19	0
	HEXLITE	31	254	OK	1	25	0
Vegas	DLVHEX	185	987	OK	2	38	0
	HEXLITE	20	191	OK	1	17	0
Marathon	DLVHEX	1,127	2,462	OK	2,543	58	29,860
	HEXLITE	28	199	OK	777	23	21,322

for solver settings. Time limits and reported times are CPU times. To make the comparison fair, we used the same PYTHON plugin for DLVHEX and HEXLITE (the plugin API of HEXLITE is compatible with the one of DLVHEX, including the API for learning nogoods in external computations). We used the latest version of DLVHEX¹¹ and the latest version of HEXLITE.¹²

5.4 Results

We will write TO (resp., MO) to indicate that the time (resp., memory) limit was exceeded.

Figure 2 shows cactus plots of the experiments on the abduction benchmark. For evaluation time (instantiation and search), for SK/P¹, HEXLITE solves each instance within at most 14s while DLVHEX fails to solve 20 instances (8 times TO and 12 times MO). For SK/G¹, HEXLITE fails to solve 32 instances because of TO and DLVHEX fails to solve 72 instances (46 times TO and 26 times MO). HEXLITE solves all instances that DLVHEX manages to solve within the timeout. The corresponding plots of external computation times show, that DLVHEX calls the external computations more often than HEXLITE to solve the same instances. The plots of memory consumption show, that memory consumption of DLVHEX rises steeper than the one of HEXLITE. Overall, HEXLITE shows a much better performance than DLVHEX on the abduction benchmark, and the difference between the solvers is more striking for SK/G¹ variant.

Table 1 shows results of running the RDF benchmark using DLVHEX and HEXLITE. Interestingly, the only timeout happens for Import where the whole set of RDF triples is represented in the answer set. DLVHEX cannot compute the result within 30 min. For Vegas, where many triples can be ignored, both DLVHEX and HEXLITE are more efficient than for Import, however DLVHEX requires 3 min while HEXLITE requires 20s to perform this (deterministic) computation. The search/optimization problem Marathon shows a big gap between DLVHEX and

¹¹ We used git hash 5a1ee06d from [git@github.com:hexhex/core.git](https://github.com/hexhex/core.git) because the stable version 2.5.0 performed significantly worse.

¹² Git hash d0e7896eb from [git@github.com:hexhex/hexlite.git](https://github.com/hexhex/hexlite.git).

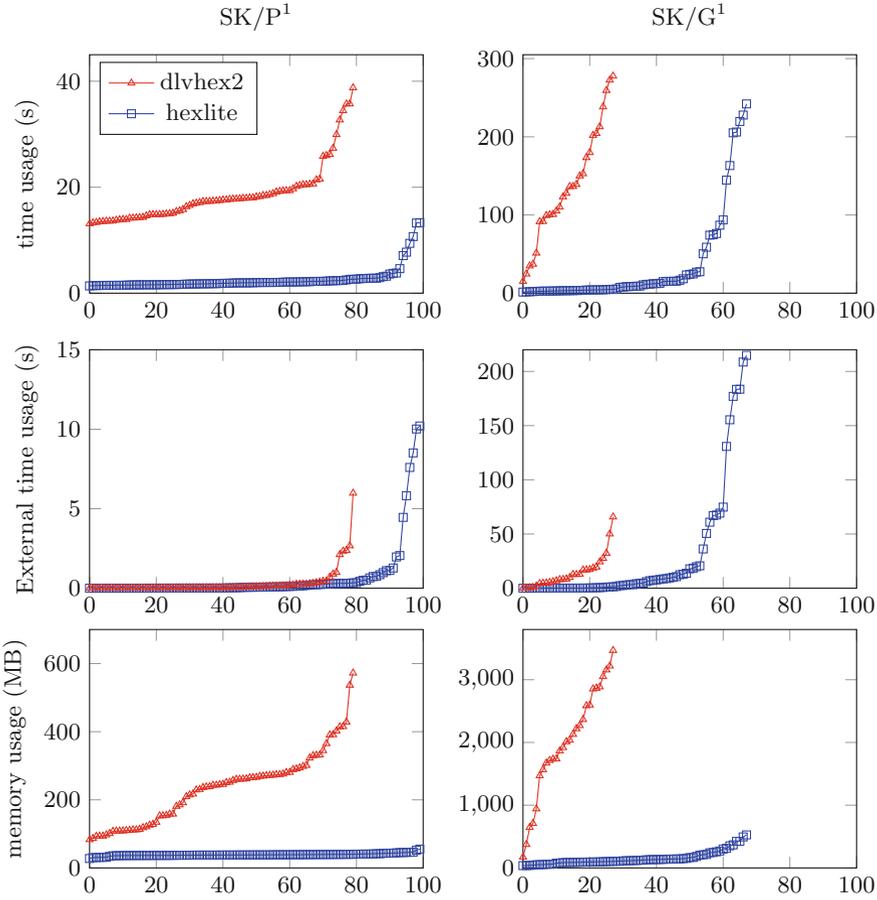


Fig. 2. Abduction benchmark: cactus plots for evaluation time, external computation time, and memory usage with two different value invention limits (SK/P¹ and SK/G¹).

HEXLITE, both in terms of time and space. In particular, HEXLITE performs fewer external atom calls and requires learning of fewer nogoods from the external computation in order to find an optimal answer set and prove its optimality. This can be explained by the lightweight usage of the CLINGO optimization feature in HEXLITE, while DLVHEX performs a lot of bookkeeping and additional computation to perform optimization.

6 Discussion and Conclusion

Our experiments show, that HEXLITE has better performance than DLVHEX in two real-world applications of HEX which are both in the Pragmatic HEX Fragment. Clearly, not all HEX programs are in PHF and some problems will be

difficult to convert into PHF. Nevertheless, wherever HEXLITE can be applied it can be a faster alternative to DLVHEX that is also easier to install as it is fully implemented in PYTHON. In cases where the program contains only external atoms of type (S) and there is a loop over at least one such external atom in the program, the FLP check becomes mandatory and the DLVHEX system will most likely have a better performance than HEXLITE because of its more advanced method for performing the FLP check [12]. If only external atoms of type (G) exist, HEXLITE will most likely perform better than DLVHEX.

There are several reasons for the better performance of HEXLITE compared with DLVHEX. Firstly, HEXLITE can remove many external atoms during grounding and they will not even be seen by the solver, while DLVHEX creates guesses for all external atoms, even those that do not depend on the interpretation, and needs to verify their truth during the solving. This eliminates a lot of potential for backend solver preprocessing in DLVHEX. Secondly, HEXLITE passes all optimization tasks to the CLINGO solver backend and just checks during propagation whether the current assignment is in itself consistent (all interpretations are checked against external atom semantics, non-partial interpretations are additionally checked against the FLP property). Opposed to that, DLVHEX performs a lot of internal bookkeeping related to optimization and maintains its own representation of the interpretation and its own cost representation for answer sets, which causes significant memory and computation overhead.

Performance issues of DLVHEX were the original reason that the tool for cost-based abduction described in [30] is formalized using the HEX formalism but implemented using a custom reasoner based on CLINGO. With HEXLITE, this implementation could be realized with less effort and without a dedicated algorithm just by using HEXLITE and implementations of two external computations.

HEXLITE uses many aspects of the CLINGO API, however it does not use CLINGO Externals [23] which are truth values given ‘from the outside’ to the solving process. Instead, the HEX notion of external computation is in a tight interaction with the answer set candidate and the results of other external computations, and during HEX evaluation, truth values of externals might be reconsidered before finding an answer set. What we do use in the FLP checker are CLASP ‘assumptions’ to communicate the answer set candidate to the FLP checker (which uses a single rewriting of the ground program for checking the FLP property of all potential answer sets).

Related to this work is a study on lazy instantiation of constraints and an alternative usage of propagators instead of constraints [7] using the WASP solver [1, 2]. In the future, WASP could be integrated into HEXLITE as an alternative to the CLINGO backend, the PYTHON API of WASP could be used to deal with external atoms of type (S). The new DLV grounder and its externals [6] could be integrated to handle external atoms of type (G).

We are glad that a study on Inductive Logic Programming [25] was successfully realized using HEXLITE. A (pragmatic) fragment of the ACTHEX extension of HEX [3, 22] based on HEXLITE is available in the HEXLITE Git repository.

Acknowledgements. We are grateful to Stefano Germano, Tobias Kaminski, Christoph Redl, Antonius Weinzierl and the anonymous reviewers for feedback about the HEXLITE system and this manuscript. This work has received funding from the Austrian Federal Ministry of Transport, Innovation and Technology (BMVIT) under grant agreement 861263 (DynaCon), and from the European Union’s Horizon 2020 research and innovation programme under grant agreement 825619 (AI4EU).

References

1. Alviano, M., Dodaro, C., Faber, W., Leone, N., Ricca, F.: WASP: a native ASP solver based on constraint learning. In: Cabalar, P., Son, T.C. (eds.) LPNMR 2013. LNCS (LNAI), vol. 8148, pp. 54–66. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40564-8_6
2. Alviano, M., Dodaro, C., Leone, N., Ricca, F.: Advances in WASP. In: Calimeri, F., Ianni, G., Truszczyński, M. (eds.) LPNMR 2015. LNCS (LNAI), vol. 9345, pp. 40–54. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-23264-5_5
3. Basol, S., Erdem, O., Fink, M., Ianni, G.: HEX programs with action atoms. In: LIPIcs-Leibniz International Proceedings in Informatics, vol. 7. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik (2010)
4. Bögl, M., Eiter, T., Fink, M., Schüller, P.: The MCS-IE system for explaining inconsistency in multi-context systems. In: Janhunen, T., Niemelä, I. (eds.) JELIA 2010. LNCS (LNAI), vol. 6341, pp. 356–359. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-15675-5_31
5. Cali, A., Gottlob, G., Kifer, M.: Taming the infinite chase: query answering under expressive relational constraints. *J. Artif. Intell. Res.* **48**, 115–174 (2013)
6. Calimeri, F., Fuscà, D., Perri, S., Zangari, J.: External computations and interoperability in the new DLV grounder. In: Esposito, F., Basili, R., Ferilli, S., Lisi, F. (eds.) AI*IA 2017. LNCS, vol. 10640, pp. 172–185. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-70169-1_13
7. Cuteri, B., Dodaro, C., Ricca, F., Schüller, P.: Constraints, lazy constraints, or propagators in ASP solving: an empirical analysis. *Theory Pract. Log. Program.* **17**(5–6), 780–799 (2017)
8. Eiter, T., Fink, M., Ianni, G., Krennwallner, T., Redl, C., Schüller, P.: A model building framework for answer set programming with external computations. *Theory Pract. Log. Program.* **16**(4), 418–464 (2016)
9. Eiter, T., Fink, M., Krennwallner, T., Redl, C.: HEX-programs with existential quantification. In: Workshop on Logic Programming, pp. 99–117 (2013)
10. Eiter, T., Fink, M., Krennwallner, T., Redl, C.: Liberal safety for answer set programs with external sources. In: AAAI Conference on Artificial Intelligence, pp. 267–275 (2013)
11. Eiter, T., Fink, M., Krennwallner, T., Redl, C.: Domain expansion for ASP-programs with external sources. *Artif. Intell.* **233**, 84–121 (2016)
12. Eiter, T., Fink, M., Krennwallner, T., Redl, C., Schüller, P.: Efficient HEX-program evaluation based on unfounded sets. *J. Artif. Intell. Res.* **49**, 269–321 (2014)
13. Eiter, T., et al.: The DLVHEX system. *KI - Künstliche Intelligenz* **32**(2), 187–189 (2018). <http://www.kr.tuwien.ac.at/research/systems/dlvhex/>
14. Eiter, T., Ianni, G., Lukasiewicz, T., Schindlauer, R., Tompits, H.: Combining answer set programming with description logics for the semantic web. *Artif. Intell.* **172**(12–13), 1495–1539 (2008)

15. Eiter, T., Ianni, G., Schindlauer, R., Tompits, H.: A uniform integration of higher-order reasoning and external evaluations in answer-set programming. In: International Joint Conference on Artificial Intelligence (IJCAI), pp. 90–96 (2005)
16. Eiter, T., Ianni, G., Schindlauer, R., Tompits, H.: Effective integration of declarative rules with external evaluations for semantic-web reasoning. In: Sure, Y., Domingue, J. (eds.) ESWC 2006. LNCS, vol. 4011, pp. 273–287. Springer, Heidelberg (2006). https://doi.org/10.1007/11762256_22
17. Eiter, T., Kaminski, T.: Exploiting contextual knowledge for hybrid classification of visual objects. In: Michael, L., Kakas, A. (eds.) JELIA 2016. LNCS (LNAI), vol. 10021, pp. 223–239. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-48758-8_15
18. Eiter, T., Kaminski, T., Redl, C., Schüller, P., Weinzierl, A.: Answer set programming with external source access. In: Ianni, G., et al. (eds.) Reasoning Web 2017. LNCS, vol. 10370, pp. 204–275. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-61033-7_7
19. Eiter, T., Kaminski, T., Redl, C., Weinzierl, A.: Exploiting partial assignments for efficient evaluation of answer set programs with external source access. *J. Artif. Intell. Res.* **62**, 665–727 (2018)
20. Erdem, E., Patoglu, V., Schüller, P.: A systematic analysis of levels of integration between low-level reasoning and task planning. *AI Commun.* **29**(2), 319–349 (2016)
21. Faber, W., Pfeifer, G., Leone, N.: Semantics and complexity of recursive aggregates in answer set programming. *Artif. Intell.* **175**(1), 278–298 (2011)
22. Fink, M., Germano, S., Ianni, G., Redl, C., Schüller, P.: ActHEX: implementing HEX programs with action atoms. In: Cabalar, P., Son, T.C. (eds.) LPNMR 2013. LNCS (LNAI), vol. 8148, pp. 317–322. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40564-8_31
23. Gebser, M., Kaminski, R., Kaufmann, B., Schaub, T.: Multi-shot ASP solving with clingo. *Theory Pract. Log. Program.* **19**, 1–56 (2018)
24. Gelfond, M., Lifschitz, V.: Classical negation in logic programs and deductive databases. *New Gener. Comput.* **9**, 365–385 (1991)
25. Kaminski, T., Eiter, T., Inoue, K.: Exploiting answer set programming with external sources for meta-interpretive learning. *Theory Pract. Log. Program.* **18**(3–4), 571–588 (2018)
26. Lassila, O., Swick, R.: Resource description framework (RDF) model and syntax specification (1999). www.w3.org/TR/1999/REC-rdf-syntax-19990222
27. Lierler, Y.: Relating constraint answer set programming languages and algorithms. *Artif. Intell.* **207**, 1–22 (2014)
28. Ng, H.T., Mooney, R.J.: Abductive plan recognition and diagnosis: a comprehensive empirical evaluation. In: Knowledge Representation and Reasoning (KR), pp. 499–508 (1992)
29. Rosis, A.D., Eiter, T., Redl, C., Ricca, F.: Constraint answer set programming based on HEX-programs. In: Proceedings of Workshop on Answer Set Programming and Other Computing Paradigms (ASPOCP) (2015)
30. Schüller, P.: Modeling variations of first-order Horn abduction in answer set programming. *Fundamenta Informaticae* **149**(1–2), 159–207 (2016)
31. Softic, S.: Conference Linked Data (COLINDA), version 1.0, last updated 30 July 2016. 149020 triples. <http://www.colinda.org>