

# *Precomputing Datalog Evaluation Plans in Large-Scale Scenarios*

ALESSIO FIORENTINO, NICOLA LEONE, MARCO MANNA, SIMONA PERRI  
and JESSICA ZANGARI

*Department of Mathematics and Computer Science, University of Calabria, Rende, Italy*  
(e-mail: [lastname@mat.unical.it](mailto:lastname@mat.unical.it)) - <https://www.mat.unical.it>

*submitted 31 July 2019; accepted 31 July 2019*

---

## Abstract

With the more and more growing demand for semantic Web services over large databases, an efficient evaluation of Datalog queries is arousing a renewed interest among researchers and industry experts. In this scenario, to reduce memory consumption and possibly optimize execution times, the paper proposes novel techniques to determine an optimal indexing schema for the underlying database together with suitable body-orderings for the Datalog rules. The new approach is compared with the standard execution plans implemented in DLV over widely used ontological benchmarks. The results confirm that the memory usage can be significantly reduced without paying any cost in efficiency.

**KEYWORDS:** Datalog, Query Answering, Ontologies, Query-plan, Data Indexing

---

## 1 Introduction

Ontological reasoning services represent fundamental features in the development of the Semantic Web. Among them, scientists are focusing their attention on the so-called *ontology-based query answering* (OBQA), where a Boolean query has to be evaluated against a *logical theory* (*knowledge base*) consisting of an extensional *database* paired with an *ontology* (Cali et al. 2009; Ortiz 2013; Amendola et al. 2018). A number of effective practical approaches proposed in the literature rewrite the query and the ontology into an equivalent Datalog program (Carral et al. 2018; Eiter et al. 2012; Kontchakov et al. 2011; Stefanoni et al. 2012; Xiao et al. 2018).

With the more and more growing availability of large databases, however, an efficient yet memory-saving evaluation of Datalog queries is arousing a renewed interest among researchers and industry experts. Typically, classical Datalog reasoners adopt sophisticated internal policies to speed-up the computation trying to limit the memory consumption. However, when the amount of data exceeds a certain size, these policies may result inadequate. This happens, for instance, for the full-fledged Datalog system I-DLV (Calimeri et al. 2017; Calimeri et al. 2019) — originally conceived as grounding engine in DLV2 (Alviano et al. 2017). Recently, to cope with large-scale scenarios, I-DLV has been further optimized and partially re-engineered by implementing novel techniques and heuristics to reduce memory consumption and possibly optimize execution times. This process gave rise to even two branches of the system called DLV2-SERVER (Leone et al. 2019) and

OWL2DLV (Allocca et al. 2019). In this paper, we present and evaluate one of the key approaches that is at the basis of the aforementioned improvements: the precomputation via Answer Set Programming (ASP) (Gelfond and Lifschitz 1991a) of an “evaluation plan” for a given Datalog program.

To understand the principles underlying the new technique, let us first recall that I-DLV historically uses strategies for join orderings and indexing that are applied rule-by-rule at runtime and that are based on local statistics over data that become available during the computation. As a result, for databases up to a few millions tuples, these strategies ensure fast evaluation at the expense of a reasonable amount of extra memory. Conversely, for databases with billions of tuples, both the time and the space used for implementing these strategies are too high. To regain usability, the idea is to precompute a global indexing schema for the underlying database associated with suitable body-orderings for all the program rules. On the one hand, this approach is less informed since, being implemented as a preprocessing phase, it cannot rely on any relevant information known during the computation, possibly leading to worsening in time. On the other hand, this allows to save both the time and space needed for computing/storing this information. Moreover, the global view on the program allows for a more parsimonious choice of the indices. To make up for the lack of local statistics, our approach is based on the natural assumption that, when dealing with very large databases, some information and statistics about the user domain are known in advance since they do not vary as fast as the actual data. This is the case, for example, for primary keys, foreign keys, small relations or selectivity of attributes. Our contribution can be therefore summarized as follows:

- Given a Datalog program  $\mathcal{P}$ , a database  $\mathcal{D}$  and some domain properties, we define the notion of *evaluation plan*, which consists of an indexing schema for  $\mathcal{D}$  together with a suitable body-ordering for each rule of  $\mathcal{P}$ . Moreover, to target “optimal” plans among all admissible ones, we identify a number of additional options, the combination of which induces different preference orderings among all plans.
- We encode the problem of finding an optimal Datalog evaluation plan in ASP, by making use of choice-rules, strong constraints, weak constraints, aggregates and negation.
- We implement optimal plans by adding annotations (Calimeri et al. 2017) to the original Datalog program  $\mathcal{P}$ . The annotated program will be the actual input for I-DLV. In this way, I-DLV execution is forced to follow the plan without the need for any internal change to the system. Nonetheless, optimal plans are sufficiently general to be implemented natively also in different Datalog engines that do not benefit from features like annotations.
- We design a well-behaved setting in the context of ontological reasoning with the aim of minimizing the memory consumption without paying in efficiency.
- We conduct an experimental evaluation over popular ontological benchmarks widely used for testing both capabilities and performance of OBQA systems. In particular, we compare performance in terms of time and memory usage of DLV when the classical computation is performed, and when the computation is driven by the planner. The results confirm that our plans improve the computation with a general gain in both time and space.

As a final remark, in case of reasoners with a server-like behavior, such as DLV2-SERVER and OWL2DLV, evaluation plans play an extremely important role, and the advantage of precomputing an evaluation plan is even more evident. Indeed, when the ontology is known in advance, it is possible to determine “offline” the optimal plan, and therefore further improve the reasoning phase with respect to both time and memory.

The present work builds on top of the extended abstract presented at JELIA 2019 (Allocca et al. 2019). In particular, apart from providing a comprehensive description of the approach, in this paper we also enrich the planner from both the formal and practical side. Indeed, we improve the notion of admissible ordering, we introduce and implement the notion of preferences, we introduce extra preferences, we identify and exploit extra info in the ontological context, and we enriched the experimental evaluation.

In the following, after recalling ASP syntax and semantics, we formalize the concept of evaluation plan for Datalog and we illustrate the modelling of such plans via ASP; eventually, we report about our experiments before drawing some conclusions.

## 2 Preliminaries

The standard input language for ASP systems is referred to as ASP-Core-2 (Calimeri et al. 2012). For the sake of simplicity, we focus next on the basic aspects of the language; for a complete reference to the ASP-Core-2 standard, and further details about advanced ASP features, we refer the reader to (Calimeri et al. 2012) and the vast literature.

A *term* is either a constant or a variable. If  $t_1, \dots, t_k$  are terms and  $p$  is a *predicate symbol* of arity  $k$ , then  $p(t_1, \dots, t_k)$  is an *atom*. A *literal*  $l$  is of the form  $a$  or  $\text{not } a$ , where  $a$  is an atom; in the former case  $l$  is *positive*, otherwise *negative*. A *rule*  $r$  is of the form  $\alpha_1 \mid \dots \mid \alpha_k \text{ :- } \beta_1, \dots, \beta_n, \text{not } \beta_{n+1}, \dots, \text{not } \beta_m$ . where  $m \geq 0$ ,  $k \geq 0$ ;  $\alpha_1, \dots, \alpha_k$  and  $\beta_1, \dots, \beta_m$  are atoms. We define  $H(r) = \{\alpha_1, \dots, \alpha_k\}$  (the *head* of  $r$ ) and  $B(r) = B^+(r) \cup B^-(r)$  (the *body* of  $r$ ), where  $B^+(r) = \{\beta_1, \dots, \beta_n\}$  (the *positive body*) and  $B^-(r) = \{\text{not } \beta_{n+1}, \dots, \text{not } \beta_m\}$  (the *negative body*). If  $H(r) = \emptyset$  then  $r$  is a (*strong*) *constraint*; if  $B(r) = \emptyset$  and  $|H(r)| = 1$  then  $r$  is a *fact*. A rule  $r$  is safe if each variable of  $r$  has an occurrence in  $B^+(r)$ <sup>1</sup>. An ASP program is a finite set  $P$  of safe rules. A program (a rule, a literal) is *ground* if it contains no variables. In the following, a Datalog program is referred to as a finite set  $P$  of safe rules stratified with respect to negation and without disjunction in the heads. A program (a rule, a literal) is *ground* if it contains no variables. A predicate is defined by a rule  $r$  if it occurs in  $H(r)$ . A predicate defined only by facts is an *EDB* predicate, the remaining are *IDB* predicates. The set of all facts in  $P$  is denoted by  $Facts(P)$ ; the set of instances of all *EDB* predicates in  $P$  is denoted by  $EDB(P)$ .

Given a program  $P$ , the *Herbrand universe* of  $P$ , denoted by  $U_P$ , consists of all ground terms that can be built combining constants and function symbols appearing in  $P$ . The *Herbrand base* of  $P$ , denoted by  $B_P$ , is the set of all ground atoms obtainable from the atoms of  $P$  by replacing variables with elements from  $U_P$ . A *substitution* for a rule  $r \in P$  is a mapping from the set of variables of  $r$  to the set  $U_P$  of ground terms. A *ground instance* of a rule  $r$  is obtained applying a substitution to  $r$ . The *full instantiation*  $Ground(P)$  of

<sup>1</sup> We remark that this definition of safety is specific for the syntax considered herein. For a complete definition we refer the reader to (Calimeri et al. 2012).

$P$  is defined as the set of all ground instances of its rules over  $U_P$ . An *interpretation*  $I$  for  $P$  is a subset of  $B_P$ . A positive literal  $a$  (resp., a negative literal *not*  $a$ ) is true w.r.t.  $I$  if  $a \in I$  (resp.,  $a \notin I$ ); it is false otherwise. Given a ground rule  $r$ , we say that  $r$  is satisfied w.r.t.  $I$  if some atom appearing in  $H(r)$  is true w.r.t.  $I$  or some literal appearing in  $B(r)$  is false w.r.t.  $I$ . Given a program  $P$ , we say that  $I$  is a *model* of  $P$ , iff all rules in  $Ground(P)$  are satisfied w.r.t.  $I$ . A model  $M$  is *minimal* if there is no model  $N$  for  $P$  such that  $N \subset M$ . The *Gelfond-Lifschitz reduct* (Gelfond and Lifschitz 1991b) of  $P$ , w.r.t. an interpretation  $I$ , is the positive ground program  $P^I$  obtained from  $Ground(P)$  by: (i) deleting all rules having a negative literal false w.r.t.  $I$ ; (ii) deleting all negative literals from the remaining rules.  $I \subseteq B_P$  is an *answer set* for a program  $P$  iff  $I$  is a minimal model for  $P^I$ . The set of all answer sets for  $P$  is denoted by  $AS(P)$ .

### 3 Datalog Evaluation Plans

In its default computational process, for optimizing the evaluation of each rule, I-DLV determines body orderings and indices on demand, according to strategies taking into account only local information for the rule at hand. In more detail, before instantiating some Datalog rule, I-DLV reorders the body literals on the basis of some join-ordering heuristics (Calimeri et al. 2017); then, according to the chosen ordering, it determines and creates needed indices. However, when memory consumption must be limited, an approach based on a global view over all rules, allowing for a more parsimonious creation of indices, is preferable.

In this section, we describe our approach for computing *optimal evaluation plans* for a set  $\mathcal{P}$  of positive Datalog rules to be evaluated over an extensional database  $\mathcal{D}$ . We define an evaluation plan of  $\mathcal{P}$  as an indexing schema over  $\mathcal{P} \cup \mathcal{D}$  together with a suitable body-ordering for each rule of  $\mathcal{P}$ . An indexing schema consists of the set of indices adopted to instantiate all rules in  $\mathcal{P}$  over  $\mathcal{D}$ . Our approach makes use of an ASP program for computing an *optimal* evaluation plan  $\mathcal{E}$  of  $\mathcal{P}$  in a preprocessing phase; then  $\mathcal{P}$  is annotated with directions that force DLV computation to follow  $\mathcal{E}$  when evaluating  $\mathcal{P}$ .

In the following, after a formal definition of an *evaluation plan*, we introduce the notion of *strategy* and then we specify the concept of *optimal* evaluation plan w.r.t. a certain strategy.

#### 3.1 Admissible Plans

Let  $\mathcal{P}$  be a set of positive Datalog rules with non-empty body and let  $\mathcal{D}$  be a database, i.e. a set of facts. We indicate with  $\text{pred}(\mathcal{P} \cup \mathcal{D})$  the set of all predicates occurring in  $\mathcal{P} \cup \mathcal{D}$  and with  $\text{rel}(p)$  the set  $\{\alpha \in \mathcal{D} : \text{pred}(\alpha) = p\}$  of the elements of  $\mathcal{D}$  sharing the predicate name  $p$ . We write  $p[i]$  to indicate the  $i$ -th argument of the predicate  $p$ . In the following, after formalizing the standard notions of *ordering* of a rule and *indexing schema* of a Datalog program, we introduce the novel notion of *evaluation plan* along with some preliminary definitions.

##### Definition 3.1

Let  $r$  be a rule in  $\mathcal{P}$  and  $B(r)$  be the set of the atoms appearing in the body of  $r$ . Let  $F_a$  be a (possibly empty) subset of atoms in  $B(r)$  and  $F_p$  be a subset of  $\{1, \dots, |B(r)|\}$ . A *position assignment* on  $r$  is a one-to-one map  $p_r : F_a \rightarrow F_p$ . A pair  $(\alpha, p)$  such that

$p_r(\alpha) = p$  is called a *fixed position* w.r.t.  $p_r$ . An *ordering* on  $r$  is a bijective function  $pos(r, \cdot) : B(r) \rightarrow \{1, \dots, |B(r)|\}$ . Having fixed a position assignment  $p_r$  on  $r$ , we define a  $p_r$ -*ordering* on  $r$  as an ordering on  $r$  such that  $pos(r, \alpha) = p_r(\alpha)$  for each  $\alpha \in F_a$ .

The definition above presents a body ordering as a rearrangement of the literals in the body, but notably, allows for having a certain number of atoms in the body in some fixed positions. This is because, according to the knowledge of the domain at hand, if one is aware that a particular choice for the orderings is convenient, the planner can be driven so that only plans complying with this choice are identified.

*Definition 3.2*

Let  $U := \{p[i] : p \in \text{pred}(\mathcal{P} \cup \mathcal{D}), 1 \leq i \leq a(p)\}$ , where  $a(p)$  represents the arity of the predicate  $p$ . An *indexing schema*  $\mathcal{S}$  over  $\mathcal{P} \cup \mathcal{D}$  is a subset of  $U$ . Given a subset  $I \subseteq U$ , we say that  $\mathcal{S}$  fixes  $I$  if  $I \subseteq \mathcal{S}$ .

Intuitively, an indexing schema is a subset of the arguments of all predicates in  $\text{pred}(\mathcal{P} \cup \mathcal{D})$ . Furthermore, similarly to the definition of ordering that may allow for fixed positions, we give the possibility to fix also a set of indices.

*Example 3.1*

As a running example in this section, consider the following positive Datalog rule  $r$  :

$$h(X, Z, W) :- a(X, Z), b(V, W), c(Z), d(V), e(Y, Z).$$

Let's consider the position assignment  $p_r$  which fixes the atom  $b(V, W)$  in first position. A possible  $p_r$ -ordering may be:

$$pos(r, a(X, Z)) = 3, pos(r, b(V, W)) = 1, pos(r, c(Z)) = 5, \\ pos(r, d(V)) = 2, pos(r, e(Y, Z)) = 4.$$

By means of such ordering the body atoms of  $r$  are rearranged as follows:

$$h(X, Z, W) :- b(V, W), d(V), a(X, Z), e(Y, Z), c(Z).$$

The set  $\mathcal{S} := \{a[2], c[1], d[1], e[2]\}$  is an example of indexing schema over the predicates appearing in  $r$ .

With a rule  $r \in \mathcal{P}$  we can associate a hypergraph  $H(r) = (V, E)$  whose vertex set  $V$  is the set of all terms appearing in  $B(r)$  and the edges in  $E$  are the term sets of each atom in  $B(r)$ . Given a rule  $r$  of  $\mathcal{P}$ , a *connected component* of  $r$  is a set of atoms in  $B(r)$  that define a connected component in  $H(r)$ .

*Example 3.2*

Let  $r$  be the rule of Example 3.1. The hypergraph  $H(r)$  associated to  $r$  has  $V = \{X, Y, Z, V, W\}$  and  $E = \{\{X, Z\}, \{V, W\}, \{Z\}, \{V\}, \{Y, Z\}\}$ . The connected components of  $r$  are  $C_1 = \{a(X, Z), c(Z), e(Y, Z)\}$  and  $C_2 = \{b(V, W), d(V)\}$ .

Let us introduce now the notions of *separation* between two connected components and *well ordering* of a component of a rule.

*Definition 3.3*

Let  $r$  be a rule of  $\mathcal{P}$  and  $pos(r, \cdot)$  be an ordering on  $r$ . Two connected components  $C_1$  and  $C_2$  of  $r$  are *separated* w.r.t.  $pos(r, \cdot)$  if  $\max\{pos(r, \alpha) : \alpha \in C_1\} < \min\{pos(r, \beta) : \beta \in C_2\}$  or vice versa.

An argument of an atom appearing in the body of a rule  $r$ , is said to be *bound*, w.r.t. an ordering on  $r$ , if it is either a constant or a variable appearing in a previous atom, and is said to be *indexBound*, w.r.t. an ordering on  $r$  and an indexing schema  $\mathcal{S}$ , if it is bound and it belongs to the schema  $\mathcal{S}$ . The definition below provides the notion of *well ordering* of a connected component in a rule.

*Definition 3.4*

Let  $r$  be a rule of  $\mathcal{P}$ ,  $\mathcal{S}$  be an indexing schema and  $pos(r, \cdot)$  be an ordering on  $r$ . A connected component  $C$  of  $r$  is *well-ordered w.r.t.  $\mathcal{S}$  and  $pos(r, \cdot)$*  if, assuming  $m = \min\{pos(r, \alpha) : \alpha \in C\}$ , for each  $\beta \in C$  with  $pos(r, \beta) = j$  and  $j > m$ , it holds that: (i)  $\beta$  has at least an argument which is indexBound, and (ii) either all the arguments of  $\beta$  are bound or there is no other atom in a later position (in the same component) that, placed in place of  $\beta$ , would have all the arguments bound.

*Example 3.3*

Let's consider the rule  $r$  with the ordering  $pos(r, \cdot)$  and the indexing schema  $\mathcal{S}$  as in our running example. The two connected components of  $r$  are clearly separated w.r.t.  $pos(r, \cdot)$ . The indexBound arguments w.r.t.  $pos(r, \cdot)$  and  $\mathcal{S}$  are  $c[1]$ ,  $d[1]$  and  $e[2]$ . It can be easily seen that the connected component  $C_2$  is well-ordered w.r.t.  $\mathcal{S}$  and  $pos(r, \cdot)$ . The same cannot be said for the component  $C_1$ ; in fact, not all the arguments of the atom  $e(Y, Z)$  are bound and the atom  $c(Z)$ , positioned in place of  $e(Y, Z)$ , would have all the arguments bound.

The notion of separation among connected components is needed for identifying, within rule bodies, clusters of literals that do not share variables. The idea is that the ordering computed by the planner should keep separated these clusters in order to avoid, as much as possible, the computation of Cartesian products during the instantiation; at the same time, literals within the clusters are properly rearranged in order to comply with the selected indexing schema, thus avoiding the creation of further indices.

Next, we provide the *admissibility* property which, in turn, characterizes the evaluation plans.

*Definition 3.5*

Given a rule  $r \in \mathcal{P}$  and an indexing schema  $\mathcal{S}$ , we say that an ordering  $pos(r, \cdot)$  is *admissible* w.r.t.  $\mathcal{S}$  if the connected components of  $r$  are mutually separated (w.r.t.  $pos(r, \cdot)$ ) and well-ordered (w.r.t.  $pos(r, \cdot)$  and  $\mathcal{S}$ ).

We define below an evaluation plan for a Datalog program.

*Definition 3.6*

Let (i)  $\{p_r; r \in \mathcal{P}\}$  be a given set of position assignments, and (ii)  $I$  be a given subset of  $\{p[i] : p \in \text{pred}(\mathcal{P} \cup \mathcal{D}), 1 \leq i \leq a(p)\}$ . An *evaluation plan*  $\mathcal{E}$  of  $\mathcal{P}$  consists of an indexing schema  $\mathcal{S}$  that fixes  $I$  together with a  $p_r$ -orderings for each  $r \in \mathcal{P}$  being admissible w.r.t.  $\mathcal{S}$ . We say that  $\mathcal{P}$  enjoys an *efficient evaluation* if it is associated to an evaluation plan.

*Example 3.4*

In our running example, the ordering  $ord(r, \cdot)$  is not admissible w.r.t. the schema  $\mathcal{S}$ . Thus,  $\mathcal{S}$  and  $ord(r, \cdot)$  do not represent an evaluation plan of the program  $\mathcal{P} = \{r\}$ . However,

an evaluation plan of  $\mathcal{P}$  could be obtained by exchanging the assignments of the atoms  $c(Z)$  and  $e(Y, Z)$  in  $ord(r, \cdot)$ . It would be appropriate to note that, in the latter case, we would obtain a further evaluation plan by excluding the argument  $a[2]$  from the indexing schema (and thus saving an index). Starting from this consideration, we introduce the concepts of preference and evaluation strategy in the next section.

### 3.2 Preferences

Let  $\mathcal{P}$  be a positive Datalog program,  $E_{\mathcal{P}}$  be the set of all the evaluation plans of  $\mathcal{P}$  and  $w : E_{\mathcal{P}} \rightarrow \mathbb{N}$  be a function that we call *cost function on  $E_{\mathcal{P}}$* . Given two evaluation plans  $\mathcal{E}_1, \mathcal{E}_2 \in E_{\mathcal{P}}$ , we say that  $\mathcal{E}_1$  is preferable to  $\mathcal{E}_2$  w.r.t. the function cost  $w$  if  $w(\mathcal{E}_1) < w(\mathcal{E}_2)$ , while we say that  $\mathcal{E}_1$  is equivalent to  $\mathcal{E}_2$  w.r.t.  $w$  if  $w(\mathcal{E}_1) = w(\mathcal{E}_2)$ . Moreover, consider a finite set  $W = \{w_1, \dots, w_n\}$  of cost functions on  $E_{\mathcal{P}}$ , we define an *evaluation strategy for  $\mathcal{P}$*  as a finite sequence  $\Sigma = (w_{\delta_1}, \dots, w_{\delta_k})$  of distinct elements in  $W$ . We say that  $\mathcal{E}_1 \in E_{\mathcal{P}}$  is preferable to  $\mathcal{E}_2 \in E_{\mathcal{P}}$  w.r.t. the strategy  $\Sigma = (w_1, \dots, w_k)$  if either:

- $\mathcal{E}_1$  is preferable to  $\mathcal{E}_2$  w.r.t.  $w_1$ , or
- there exists  $j \in \{2, \dots, k\}$  such that  $\mathcal{E}_1$  is equivalent to  $\mathcal{E}_2$  w.r.t.  $w_i$  for each  $i = 1, \dots, j - 1$ , and  $\mathcal{E}_1$  is preferable to  $\mathcal{E}_2$  w.r.t.  $w_j$ .

According to the notion of *preference* of an evaluation plan over another w.r.t. a strategy, we can now introduce the definition of “optimal” plans w.r.t. that strategy.

#### Definition 3.7

Let  $E_{\mathcal{P}}$  be the set of all the evaluation plans for a positive Datalog program  $\mathcal{P}$  and  $\Sigma$  be an evaluation strategy for  $\mathcal{P}$ . An evaluation plan  $\mathcal{E}_0$  is said to be *optimal w.r.t.  $\Sigma$*  if it is either preferable or equivalent to each  $\mathcal{E} \in E_{\mathcal{P}}$  w.r.t.  $\Sigma$ .

Intuitively, finding the optimal evaluation plans against a strategy  $\Sigma = (w_1, \dots, w_k)$  means finding those that minimize the cost function  $w_1$ , then, among these, find those that minimize the function  $w_2$ , and so on. We report next four functions used for defining our evaluation strategies. From now on when we talk about the functions  $w_1, w_2, w_3$  and  $w_4$  we will refer to the following:

- $w_1(\mathcal{E}) := \sum_{p[i] \in \mathcal{S}} c(p, i)$ , where  $c(p, i)$  is the cost of building an index over  $p[i]$  in the indexing schema  $\mathcal{S}$ . Note that we presuppose the knowledge of  $c(p, i)$  values. Such costs can be estimated via heuristics or actually computed, depending on the application domain at hand. As said in the introduction, the novel approach is based on the natural assumption that, when dealing with very large databases, some information and statistics about the user domain are known in advance since they do not vary as fast as the actual data. Apart from primary keys and foreign keys, which in OBQA are related to the ontological schema, some statistics on the data can be also taken into account. This is the case, for example, of the *estimation of the selectivity* of an attribute, which gives an indication of the average number of times that a constant (or individual) occurs in the relation in correspondence of the given attribute (note that, the special case of estimation of the selection equal to 1 indicates that the given attribute is actually a key). This value can be taken into account for estimating the size (and thus, the cost) of an index for the given attribute.

- $w_2(\mathcal{E})$  is defined as the sum of the positions of atoms involved in recursion. We prefer that atoms involved in recursion are placed as soon as possible. The extension of such atoms might considerably grow and change during computation; placing them before other atoms in the body could avoid the creation of expensive indices.
- $w_3(\mathcal{E})$  is the number of indices set on arguments that are not primary keys. In other words we prefer indices set on arguments representing primary keys.
- $w_4(\mathcal{E}) := \sum_{r \in \mathcal{P}} \sum_{\alpha \in B(r)} [\maxArity - u(\alpha, r)] * pos(r, \alpha)$ , where  $\maxArity$  represents the maximum arity of the atoms appearing in  $\mathcal{P}$  and  $u(\alpha, r)$  is the number of unbound arguments of the atom  $\alpha$  in the rule  $r$ . We prefer that atoms having large number of unbound arguments (that is, those that minimize the first factor in the above summation) are placed as soon as possible as they possibly will lead to have new completely bound atoms to be placed in successive positions.

## 4 ASP-based Implementation

In the following we describe the ASP code devised in order to compute optimal evaluation plans. For the sake of simplicity, as the program is rather long and involved, we report here only some key parts; the full ASP code is available online.<sup>2</sup>

The program is based on the classical “Guess/Check/Optimize” paradigm and combines: (i) *choice and disjunctive rules* to guess an indexing schema  $\mathcal{S}$  over  $\mathcal{P} \cup \mathcal{D}$  and, for each rule  $r$  in  $\mathcal{P}$ , an ordering  $\text{ord}(r, \cdot)$ ; (ii) *strong constraints* to guarantee, for each rule  $r$ , the admissibility of  $\text{ord}(r, \cdot)$  w.r.t.  $\mathcal{S}$ ; (iii) *weak constraints* to find out the optimal evaluation plans of  $\mathcal{P}$  w.r.t. the chosen strategy.

### 4.1 Data Model

The planner consists of an ASP program taking as input a set of facts representing  $\mathcal{P}$  and the database  $\mathcal{D}$ ; each rule of  $\mathcal{P}$  is represented by means of facts of the form:

```
rule(Rule, Description, NumberOfBodyAtoms).
headAtom(Rule, Atom, Predicate).
bodyAtom(Rule, Atom, Predicate).
sameVariable(Rule, Atom1, Arg1, Atom2, Arg2).
constant(Rule, Atom, Arg).
```

Facts over the predicate `rule` associate each rule  $r$  to an identifier and provide the number of its body atoms. Atoms in the body and in the head of each rule  $r$  are represented by `bodyAtom` and `headAtom` predicates respectively. The predicate `sameVariable` provides the common variables related to every pair of atoms appearing in  $r$ , whereas `constant` states that a constant term occurs in the argument of an atom of  $r$ . An example of the basic input concepts described above is the following:

#### Example 4.1

The following program  $\mathcal{P}$  :

```
h1(X) :- a(X, Y), b(Y).
h2(Y) :- a(Y, X).
```

<sup>2</sup> See <https://www.mat.unical.it/perri/iclp2019.zip>.



is represented by means of the facts:

```
rule(0, "h1(X) :- a(X, Y), b(Y).", 2).
headAtom(0, "h1(X)", "h1/1").
bodyAtom(0, "a(X, Y)", "a/2").
bodyAtom(0, "b(Y)", "b/1").
sameVariable(0, "h1(X)", 1, "a(X, Y)", 1).
sameVariable(0, "a(X, Y)", 2, "b(Y)", 1).

rule(1, "h2(Y) :- a(Y, X).", 1).
headAtom(1, "h2(Y)", "h2/1").
bodyAtom(1, "a(Y, X)", "a/2").
sameVariable(1, "h2(Y)", 1, "a(Y, X)", 1).
```

The database  $\mathcal{D}$  is represented by means of facts over predicate `relation`, while the costs of building indices over arguments are given by facts over predicate `index`.

```
relation(Predicate, Arity).
index(Predicate, Arg, Cost).
```

Furthermore, the planner allows for having a certain number of atoms in the body in some previously fixed positions and a set of indices fixed in  $\mathcal{S}$ . This is because, according to the knowledge of the domain at hand, if one is aware that a particular choice for the orderings and the indexing policy is convenient, the planner can be driven so that only plans complying with this choice are identified. The planner can also exploit the presence of arguments representing primary keys for predicates in  $\mathcal{P} \cup \mathcal{D}$ . Such information, if available, can be given in input to the ASP planner by means of facts of the form:

```
fixedPosition(Rule, Atom, Pos).
fixedIndex(Predicate, Arg).
key(Predicate, Arg).
```

## 4.2 Guess Part

The following choice rule (Calimeri et al. 2012) guesses a subset of the arguments of all predicates, namely an indexing schema  $\mathcal{S}$ , over  $\mathcal{P} \cup \mathcal{D}$ . Notably, the arguments to be indexed are chosen among a restricted set of arguments, called *indexable*, in order to keep the search space smaller. For instance, arguments that are not involved in joins are not indexable.

```
{setIndex(Predicate, Arg)} :- indexable(Predicate, Arg).
```

Beside this choice rule, the guess part contains also the following rule for guessing a body-ordering for each rule  $r$  in  $\mathcal{P}$ . In particular, the choice guesses a position in the body for each atom whose position has not been previously fixed (`fixedAtomRule`). Clearly, only positions not already occupied by another body atom in the same rule (`fixedPositionRule`) are guessable. The predicates `fixedAtomRule` and `fixedPositionRule` are computed according to the predicate `fixedPosition` described above.

```
{pos(Atom, Rule, Pos) : position(Pos), Pos >= 1, Pos <= Size,
 not fixedPositionRule(Rule, Pos)} = 1 :- rule(Rule, Size),
 bodyAtom(Rule, Atom, _), not fixedAtomRule(Rule, Atom).
```

### 4.3 Check Part

This part discards, by means of *strong constraints*, solutions that do not satisfy (according to the definitions in Section 3.1) the conditions to be considered admissible evaluation plans. In particular, conditions that have to be necessarily satisfied are the following:

1. The *connected components* of each rule of  $\mathcal{P}$  must be kept separate. According to the definition 3.3, this is ensured by the following constraint.

```
:- pos(Atom1,Rule,Pos1),pos(Atom2,Rule,Pos2),
   sameComponent(Rule,Atom1,Atom2),pos(Atom3,Rule,Pos3),
   not sameComponent(Rule,Atom1,Atom3),Pos1<Pos3,Pos3<Pos2.
```

2. According to the first point of the definition 3.4, to guarantee that each connected component is *well-ordered*, each atom, except those in the first position of each component, must have at least an argument *indexBound*. This condition is guaranteed by the constraint below, where predicates `firstPosition` and `indexBound` suggest, respectively, the first positions of the components in each rule, and the `indexBound` arguments of each atom in a rule.

```
:- pos(Atom,Rule,Pos),firstPosition(Rule,FirstPos),
   Pos>FirstPos,#count{Arg:indexBound(Arg,Atom,Rule)}=0.
```

3. The second condition of the definition 3.4 is modeled by means of the following constraint. Here the predicate `atomVars` indicates the number of variables occurring in every atom.

```
:- pos(Atom,Rule,Pos),not boundAtom(Atom,Rule,Pos),
   pos(Atom1,Rule,Pos1),not boundAtom(Atom1,Rule,Pos1),
   Pos1>Pos,Pos2>Pos1,boundAtom(Atom2,Rule,Pos2),
   atomVars(Atom2,Rule,N),
   #count{Arg2:sameVariable(Rule,Atom2,Arg2,Atom,_)}=N.
```

The checking part contains also an additional constraint encoding the following basic check for guaranteeing the correctness of the plans. In particular, this basic check ensures that two different atoms do not occupy the same position in any rule:

```
:- pos(Atom1,Rule,Pos),pos(Atom2,Rule,Pos),Atom1!=Atom2.
```

### 4.4 Optimize Part

Eventually, in this section we describe the part for identifying the optimal evaluation plan according to the evaluation strategy that one decides to apply. Remember that a strategy is a finite combination of cost functions. Currently, the planner is equipped with the four cost functions described in Section 3.2, each of which is represented by a specific *weak constraint*. Note that, weak constraints allow for expressing preferences possibly having different importance levels. The planner allows to fix these priority levels according to the chosen strategy by providing in input facts which indicate that the cost function  $w_N$  has priority level  $P$ .

```
priorityCostFunction(N,P).
```

For instance, suppose we want to represent the strategy  $\Sigma = (w_1, w_3, w_2)$ , then we need the following input facts to indicate that the cost function  $w_1$  has priority level 3,  $w_3$  has priority level 2 and  $w_2$  has priority level 1. Note that in this case the cost function  $w_4$  is not activated.

```

priorityCostFunction(1,3).
priorityCostFunction(3,2).
priorityCostFunction(2,1).

```

This means that the planner is customizable. Indeed, depending on the knowledge of the domain at hand, one can choose to adapt the strategy to his own needs simply by exchanging the priority levels of the cost functions among those already present in the planner, or even by integrating new cost functions (with the addition of new constraints in the encoding). In the following we illustrate the weak constraints representing the cost functions defined in Section 3.2.

1. The rule below aims to minimize index occupation. To this end, we presuppose the knowledge of the costs (or their estimation) of building indices over arguments and we represent them by facts of form `index(Predicate, Arg, Cost)`.

```

:~ setIndex(Predicate, Arg), index(Predicate, Arg, Cost),
   priorityCostFunction(1,P). [Cost@P, Predicate, Arg, Cost]

```

2. We prefer that atoms involved in recursion are placed as soon as possible. The weak constraint makes use of the auxiliary predicate `recursivePredicate` providing information about which predicates of the program are recursive. We do not report its definition for the sake of readability.

```

:~ pos(Atom, Rule, Pos), bodyAtom(Rule, Atom, Predicate),
   recursivePredicate(Predicate), priorityCostFunction(2,P).
   [Pos@P, Rule, Pos]

```

3. Indices set on arguments representing primary keys are possibly preferred:

```

:~ setIndex(Predicate, Arg), not key(Predicate, Arg),
   priorityCostFunction(3,P). [1@P, Predicate, Arg]

```

4. Atoms having large number of unbound arguments should be placed as soon as possible in the body. Also in this case we make use of an auxiliary predicate: `numBoundArgs` provides the number of bound arguments of an atom in a rule.

```

:~ numBoundArgs(Atom, Rule, Pos, B), maxArity(N),
   bodyAtom(Rule, Atom, Predicate), relation(Predicate, Arity),
   priorityCostFunction(4,P). [(N-Arity+B)*Pos@P, Rule, Pos]

```

## 5 Experimental Evaluation

Hereafter we report the results of an experimental activity carried out to assess the effectiveness of the ASP-based evaluation planner.

### 5.1 Benchmarks

Our experimental analysis relies on four benchmarks: LUBM (Lehigh University Benchmark), LUBM-LUTZ, Stock Exchange and Vicodi.

**LUBM.** It is one of the most popular ontologies for testing both capabilities and performance of OBQA systems; indeed, it has been specifically developed to facilitate the evaluation of Semantic Web reasoners in a standard and systematic way. In particular, the benchmark is intended to evaluate performance of those reasoners with respect to extensional queries over large databases that refer to a single realistic ontology. The LUBM

benchmark consists of a university domain OWL 2 ontology along with customizable and repeatable synthetic data and a set of 14 SPARQL queries<sup>3</sup>. Queries 2, 6, 9 and 14 involve constants, while the other queries are constant-free. In our experiments, the original LUBM ontology and the official 14 queries have been translated into Datalog via the CLIPPER system (Eiter et al. 2012). The official LUBM generator has been adopted to generate four databases of increasing sizes: LUBM-500, LUBM-1,000, LUBM-2,000 and LUBM-4,000, where the number associated to each database name indicates the number of universities composing it. The number of facts in the databases ranges from about 67,000,000 to about half a billion facts.

**LUBM-LUTZ.** It is a variant of LUBM designed by Lutz et al. (2013). This benchmark consists of an OWL2 ontology and 11 queries (both different from those of LUBM) along with a modified version of the LUBM official generator allowing to set the level of incompleteness in the database. As done for LUBM, the ontology and the queries have been translated into Datalog via the CLIPPER system (Eiter et al. 2012). All queries are without constants. We generated five databases of increasing sizes and having an incompleteness percentage of 10%: LUTZ-500, LUTZ-1,000, LUTZ-2,000, LUTZ-4,000 and LUTZ-8,000. Again, the number associated to each database name indicates the number of universities composing it.

**Stock Exchange and Vicodi.** These are two real world ontologies widely used in literature for the evaluation of query rewriting systems (Mora and Corcho 2013). For each of these two ontologies, we selected 5 queries featuring constants and we used the SyGENiA generator (Grau et al. 2012) to produce five databases having from 1,000 to 40,000 tuples and a number of individuals varying from 100 to 4,000. These are the maximum sizes that can be generated using SyGENiA.

## 5.2 Setting

Experiments on LUBM-LUTZ have been performed on a Dell Linux server with an Intel Xeon Gold 6140 CPU composed of 8 physical CPUs clocked at 2.30 GHz, with 297GB of RAM. Experiments on LUBM, Vicodi and Stock Exchange have been performed on a NUMA Linux machine equipped with two 2.8 GHz AMD Opteron 6320 processors and 128GB RAM. Unlimited time and memory were granted to running processes. Benchmarks and executables used for the experiments are available at <https://www.mat.unical.it/perri/iclp2019.zip>.

Two different executions have been compared: (i) a classical execution of I-DLV which, given as input the so generated encodings, chooses body orderings and indexing strategies with its default policies, and (ii) an execution driven by the planner in which I-DLV is forced to follow the precomputed evaluation plan that decided body orderings and indices in order to reduce memory consumption. These constraints have been defined via annotations, that represent specific means to express preferences over its internal computational process (Calimeri et al. 2017).

<sup>3</sup> LUBM is available at <http://swat.cse.lehigh.edu/projects/lubm/>.

### 5.3 Planner customization

In the context of OBQA, where the objective is to answer a query, the rewritten Datalog program typically benefits from the application of the so-called Magic Sets technique (Alviano et al. 2012). This produces a new equivalent program containing extra intensional predicates that could have very small extensions during the computation. These predicates, in a setting where memory consumption should be limited, could be moved towards the end of the body so that it is more likely saving space for needed indices. Hence, to instantiate our planner, we consider this additional domain information and use facts of the form `fixedPosition(Rule,Atom,Pos)` to specify it. It is worth remarking that in I-DLV this customization has an impact only in case of queries featuring constants since magic atoms are not generated for queries without constants. Furthermore, for all attributes involved in extensional relations, we provide, via facts of the form `index(Predicate,Arg,Cost)`, an estimation of the size of an index for that attribute. In particular, in our experiments, to have available this information we generate and analyze a “small” database for each benchmark.

In our experiments, we considered different planner customizations depending on the domain at hand. In particular, for LUBM, Vicodi and Stock Exchange, consisting mainly of queries featuring constants, we adopted the strategy  $\Sigma_1 = (w_2, w_4)$ . The idea underlying this choice is that, in such domains, I-DLV can benefit from the Magic Sets technique and fixing the positions of magic atoms as described above is already sufficient to drive the planner. On LUBM-LUTZ, having instead constant-free queries, we adopted the strategy  $\Sigma_2 = (w_1, w_2, w_3, w_4)$ ; indeed, since Magic Sets are not active, no fixed positions can be provided and a richer strategy is necessary for avoiding an almost blind plan computation.

### 5.4 Discussion

The results of our experiments are reported in Table 1 and in Figure 1 and 2.

Table 1 shows performance in terms of average running time and memory usage of I-DLV (with and without planner) computed over all considered databases per each benchmark query. Columns 2 and 3 refer to the classical computation, while columns 4 and




Fig. 1. Experiments on LUBM. Queries are ordered by increasing values w.r.t. the No-Planner execution.




Fig. 2. Experiments on LUBM-LUTZ. Queries are ordered by increasing values w.r.t. the No-Planner execution.

5 to the computation driven by the planner. In the 6th column, we reported the time spent to compute the optimal plan, in the 7th column, the memory saving per query computed as difference of the corresponding fields in columns 3 and 5. Similarly, the 8th column reports the time saving per query computed as difference of the corresponding fields in columns 2 and 4. The table reports also some aggregated data per benchmark. In particular, it shows information on the average/maximum saved memory/time, as well as the number of queries where an improvement in terms of saved memory (resp. saved time) has been obtained. In addition, to provide a clearer picture of the behavior of the two versions of I-DLV, we reported in Figures 1 and 2 plots of the average running time and memory usage over all considered databases for the largest benchmarks: LUBM and LUBM-LUTZ.

As it can be seen, we obtained a significant saving of memory on LUBM where, for instance, the planner allows to save 11.4 GB on query *q13* (about 40% less) w.r.t. the no-planner version, and a gain both in terms of memory and time over almost all queries. Only on query *q01* we experimented a small worsening on memory. In general, no significant increase of computation time is observable and, in several cases, the planner-driven approach leads also to improvements in terms of time. This can be explained considering that indices selected by the planner, being on the overall less memory expensive, are more efficiently computable.

Concerning LUBM-LUTZ, we first note that the benefits appear less evident. This is due to the nature of the queries in the benchmark which are constant-free and require a different (less informed) customization, as described in Section 5.3. Nonetheless, the execution of I-DLV driven by the planner performs generally better (both in time and memory) of the standard execution. On the queries *q08*, *q09* and *q11* we have a memory saving of 9-10% w.r.t. the no-planner version; moreover, we observe no worsening in memory consumption and only one case in which there is a negligible worsening in time.

As for Stock Exchange and Vicodi, although these are not data intensive domains, I-DLV can benefit by the planner as well. Indeed, worsenings in terms of memory range from 1% to 7% in a few queries which are somehow expected when measuring memory of the order of megabytes.

Further aggregated data and statistics on the results are given in Table 2. This shows, for both the tested versions of I-DLV and for each benchmark, the maximum peak of

Table 1. Experiments on LUBM, LUBM-LUTZ, Stock Exchange and Vicodi. Time is in seconds, memory is in GB for LUBM and LUBM-LUTZ and in MB for Stock Exchange and Vicodi.

Query	No Planner		Planner				
	Time	Memory	Time	Memory	Planning Time	Saved Memory	Saved Time
LUBM (Memory in GB, Time in seconds)							
q01	1949.24	17.27	1526.94	17.67	0.01	-0.4	422.3
q02	2292.53	19.22	1121.70	17.27	0.14	2.0	1170.8
q03	1643.14	17.89	1026.31	17.27	0.01	0.6	616.8
q04	2250.26	37.21	1241.08	33.28	0.3	3.9	1009.2
q05	1303.48	28.28	1087.97	17.27	0.17	11.0	215.5
q06	1485.75	23.74	1383.24	21.12	0.05	2.6	102.5
q07	1284.16	25.94	1124.31	21.23	0.18	4.7	159.9
q08	1349.36	32.47	1162.35	22.49	0.32	10.0	187.0
q09	1464.82	23.74	1326.15	21.32	0.22	2.4	138.7
q10	1271.19	24.94	1101.55	20.84	0.14	4.1	169.6
q11	1016.69	17.27	1014.10	17.27	0.01	0.0	2.6
q12	1275.24	27.13	1115.07	18.68	0.2	8.4	160.2
q13	1281.30	30.06	1129.73	18.66	0.17	11.4	151.6
q14	1122.83	17.27	1080.39	17.27	0	0.0	42.4
Saved Memory	Average: 4.34		Maximum: 11.4		Improvements: 13/14		
Saved Time	Average: 324.94		Maximum: 1170.83		Improvements: 14/14		
LUTZ (Memory in GB, Time in seconds)							
q01	353.88	8.17	328.94	8.15	0	0.0	24.9
q02	237.25	6.19	232.24	6.20	0.04	0.0	5.0
q03	268.90	6.63	274.70	6.20	0.26	0.4	-5.8
q04	249.86	6.67	240.43	6.63	0.02	0.0	9.4
q05	246.67	6.21	240.60	6.21	0.08	0.0	6.1
q06	273.83	6.88	258.81	6.63	0.28	0.2	15.0
q07	219.15	6.20	218.20	6.20	0.05	0.0	0.9
q08	293.97	7.59	284.52	6.80	0.11	0.8	9.5
q09	260.12	6.80	247.42	6.20	0.01	0.6	12.7
q10	289.85	7.19	283.63	7.22	0.04	0.0	6.2
q11	288.08	7.14	280.95	6.44	0.07	0.7	7.1
Saved Memory	Average: 0.26		Maximum: 0.79		Improvements: 11/11		
Saved Time	Average: 8.28		Maximum: 24.94		Improvements: 10/11		
Stock Exchange (Memory in MB, Time in seconds)							
q01	0.83	11.98	0.83	12.16	0	-0.2	0.0
q02	1.10	23.48	1.15	21.10	0.02	2.4	-0.1
q03	2.06	39.64	2.41	36.60	0.03	3.0	-0.4
q04	1.31	29.04	1.32	25.28	0.04	3.8	0.0
q05	2.35	46.40	2.44	41.34	0.09	5.1	-0.1
Saved Memory	Average: 2.81		Maximum: 5.06		Improvements: 4/5		
Saved Time	Average: -0.1		Maximum: 0		Improvements: 2/5		
Vicodi (Memory in MB, Time in seconds)							
q01	0.83	9.88	0.83	9.96	0.01	-0.1	0.0
q02	0.93	20.56	0.83	12.02	0.01	8.5	0.1
q03	0.82	11.38	0.78	10.72	0.02	0.7	0.0
q04	0.70	12.88	0.70	13.84	0.02	-1.0	0.0
q05	0.78	12.54	0.78	13.10	0.08	-0.6	0.0
Saved Memory	Average: 1.52		Maximum: 8.54		Improvements: 2/5		
Saved Time	Average: 0.03		Maximum: 0.11		Improvements: 5/5		

memory and the total sum of execution times computed over all databases and queries, along with the corresponding profits. In all benchmarks, the peak of memory when the

Table 2. *Statistics on LUBM, LUBM-LUTZ, Stock Exchange and Vicodi: the maximum peak of memory and the total sum of execution times computed over all databases and queries, along with the corresponding profits. Time is in seconds, memory is in MB.*

Domain	Peak of Memory			Sum of Times		
	No Planner	Planner	Profit (%)	No Planner	Planner	Profit (%)
LUBM	81439.90	72733.30	10.7%	83960.03	65763.48	21.7%
LUBM-LUTZ	21701.70	21634.20	0.3%	14907.74	14452.21	3.1%
Stock Exchange	77.40	68.90	11.0%	38.25	40.72	-6.5%
Vicodi	38.50	26.60	30.9%	20.32	19.58	3.7%

planner is used is less than the one obtained using the standard version of I-DLV. Regarding times, although we experimented a small worsening for Stock Exchange (6.5%), we observe a general improvement which is greater than 20% in our large-scale benchmark.

## 6 Conclusion

In this work we introduced an evaluation planner for Datalog programs. The planner has been conceived to be applied to ontology-based query answering contexts, where often, in case of large databases, standard approaches are not convenient/applicable due to memory consumption. It relies on an ASP program that computes the plan, intended as an indexing schema for the database together with a body-ordering for each rule in the program. The computed plan minimizes the overall cost (in term of memory consumption) of indices; moreover, the usage of the plan with the DLV system allows to further reduce memory usage since some expensive internal optimizations of DLV can be disabled. Results of the experiments conducted on popular ontological benchmarks confirm the effectiveness of the approach. Eventually, precomputing offline an evaluation plan plays an extremely important role in reasoners with a server-like behavior, since this allows for further reducing the time required by the actual computation.

## Acknowledgments

This work has been partially supported by MIUR under project “Declarative Reasoning over Streams” (CUP H24I17000080001) – PRIN 2017, by MISE under project “S2BDW” (F/050389/01-03/X32) – “Horizon2020” PON I&C2014-20, by Regione Calabria under project “DLV LargeScale” (CUP J28C17000220006) – POR Calabria 2014-20.

## References

- ALLOCCA, C., CALIMERI, F., CIVILI, C., COSTABILE, R., CUTERI, B., FIORENTINO, A., FUSCÀ, D., GERMANO, S., LABOCCETTA, G., MANNA, M., PERRI, S., REALE, K., RICCA, F., VELTRI, P., AND ZANGARI, J. 2019. Large-scale reasoning on expressive horn ontologies. In *Proceedings of Datalog 2.0*, M. Alviano and A. Pieris, Eds. CEUR Workshop Proceedings, vol. 2368. CEUR-WS.org, 10–21.
- ALLOCCA, C., COSTABILE, R., FIORENTINO, A., PERRI, S., AND ZANGARI, J. 2019. Memory-saving evaluation plans for datalog. In *JELIA*. Lecture Notes in Computer Science, vol. 11468. Springer, 453–461.



- ALVIANO, M., CALIMERI, F., DODARO, C., FUSCÀ, D., LEONE, N., PERRI, S., RICCA, F., VELTRI, P., AND ZANGARI, J. 2017. The ASP system DLV2. In *Proc. of LPNMR*. LNCS, vol. 10377. 215–221.
- ALVIANO, M., FABER, W., GRECO, G., AND LEONE, N. 2012. Magic sets for disjunctive datalog programs. *Artificial Intelligence* 187, 156–192.
- AMENDOLA, G., LEONE, N., AND MANNA, M. 2018. Finite controllability of conjunctive query answering with existential : Two steps forward. In *IJCAI*. 5189–5193.
- CALÌ, A., GOTTLÖB, G., AND LUKASIEWICZ, T. 2009. Tractable query answering over ontologies with datalog+/- . In *Proceedings of DL'09*.
- CALIMERI, F., FABER, W., GEBSER, M., IANNI, G., KAMINSKI, R., KRENNWALLNER, T., LEONE, N., RICCA, F., AND SCHAUB, T. 2012. Asp-core-2: Input language format. <https://www.mat.unical.it/aspcomp2013/files/ASP-CORE-2.03b.pdf>.
- CALIMERI, F., FUSCÀ, D., PERRI, S., AND ZANGARI, J. 2017. I-DLV: the new intelligent grounder of DLV. *Intelligenza Artificiale* 11, 1, 5–20.
- CALIMERI, F., PERRI, S., AND ZANGARI, J. 2019. Optimizing answer set computation via heuristic-based decomposition. *TPLP*, 126.
- CARRAL, D., DRAGOSTE, I., AND KRÖTZSCH, M. 2018. The combined approach to query answering in horn-alchoiq. In *KR*. AAAI Press, 339–348.
- EITER, T., ORTIZ, M., SIMKUS, M., TRAN, T., AND XIAO, G. 2012. Query rewriting for horn-shiq plus rules. In *Proc. of AAAI*.
- GELFOND, M. AND LIFSCHITZ, V. 1991a. Classical Negation in Logic Programs and Disjunctive Databases. *New Generation Computing* 9, 3/4, 365–385.
- GELFOND, M. AND LIFSCHITZ, V. 1991b. Classical negation in logic programs and disjunctive databases. *New Generation Comput.* 9, 3/4, 365–386.
- GRAU, B. C., MOTIK, B., STOILOS, G., AND HORROCKS, I. 2012. Completeness guarantees for incomplete ontology reasoners: Theory and practice. *J. Artif. Intell. Res.* 43, 419–476.
- KONTCHAKOV, R., LUTZ, C., TOMAN, D., WOLTER, F., AND ZAKHARYASCHEV, M. 2011. The combined approach to ontology-based data access. In *IJCAI*.
- LEONE, N., ALLOCCA, C., ALVIANO, M., CALIMERI, F., CIVILI, C., COSTABILE, R., FIORENTINO, A., FUSCÀ, D., GERMANO, S., LABOCCETTA, G., CUTERI, B., MANNA, M., PERRI, S., REALE, K., RICCA, F., VELTRI, P., AND ZANGARI, J. 2019. Enhancing DLV for large-scale reasoning. In *Proceedings of LPNMR*, M. Balduccini, Y. Lierler, and S. Woltran, Eds. Lecture Notes in Computer Science, vol. 11481. Springer, 312–325.
- LUTZ, C., SEYLAN, I., TOMAN, D., AND WOLTER, F. 2013. The combined approach to OBDA: taming role hierarchies using filters. In *International Semantic Web Conference (1)*. Lecture Notes in Computer Science, vol. 8218. Springer, 314–330.
- MORA, J. AND CORCHO, Ó. 2013. Towards a systematic benchmarking of ontology-based query rewriting systems. In *International Semantic Web Conference (2)*. LNCS, vol. 8219. Springer, 376–391.
- ORTIZ, M. 2013. Ontology based query answering: The story so far. In *AMW*. CEUR Workshop Proceedings, vol. 1087. CEUR-WS.org.
- STEFANONI, G., MOTIK, B., AND HORROCKS, I. 2012. Small datalog query rewritings for EL. In *DL*. CEUR Workshop Proceedings, vol. 846.
- XIAO, G., CALVANESE, D., KONTCHAKOV, R., LEMBO, D., POGGI, A., ROSATI, R., AND ZAKHARYASCHEV, M. 2018. Ontology-based data access: A survey. In *IJCAI*.