

RDFLog: It’s like Datalog for RDF

François Bry¹, Tim Furche¹, Clemens Ley², Benedikt Linse¹, and Bruno Marnette²

¹ Institute for Informatics, University of Munich,
Oettingenstraße 67, D-80538 München, Germany

² Oxford University Computing Laboratory,
Wolfson Building, Parks Road, Oxford, OX1 3QD, England

Abstract. RDF data is set apart from relational or XML data by its support of rich existential information in the form of *blank nodes*. Where in SQL databases null values are scoped over a single tuple, blank nodes in RDF can span over any number of statements and thus can be seen as existentially quantified variables.

Blank node querying is considered in most RDF query languages, but blank node construction, i.e., the introduction of new blank nodes has been mostly ignored (e.g., in Triple) or treated in a very limited form (e.g., in SPARQL). In this paper, we classify three kinds of blank nodes in RDF query languages and introduce the recursive, rule-based RDF query language RDFLog. RDFLog is the first RDF query language with full arbitrary quantifier alternation: blank nodes may occur in the scope of all, some, or none of the universal variables of a rule. RDFLog is also aware of important RDF features such as the distinction between blank nodes, literals and URIs or the RDFS vocabulary.

1 Introduction

Access to data in a machine-processable, domain-independent manner plays a central role in the future growth of the Internet. Information on legislative proceedings, census data, scientific experiments and databases, as well as the data gathered by social network applications is now accessible in form of RDF data. The Resource Description Framework (RDF) is a data format for the Web with a formal semantics that is achieving considerable popularity. Compared to relational databases, RDF is mostly distinguished by (1) a specialization to ternary statements or “triples” relating a subject, via a predicate, to an object, (2) the presence of blank nodes that allow statements where subject or object are unknown, and (3) specific semantics for a small, predefined vocabulary (RDF Schema, or RDFS) reminiscent of an object-oriented type system.

With the staggering amount of data available in RDF form on the Web, the second indispensable ingredient becomes the easy selection and processing of RDF data. For that purpose, a large number of RDF query languages (see [1] for a recent survey) have been proposed. In this paper, we add a further exemplar: RDFLog extends datalog to support the distinguishing features of RDF such as blank nodes and the logical core [2] of the RDFS vocabulary. In

RDFLog, Blank nodes can be constructed by existentially quantified variables in rule heads. RDFLog allows *full alternation* between existential and universal quantifiers in a rule. This sharply contrasts with previous approaches to rule-based query languages that either do not support blank nodes (in rule heads) at all [3,4], or only a limited form of quantifier alternation [5,6,7].

To illustrate the benefits of full quantifier alternation, imagine an information system about university courses. We distinguish three types of rules with existential quantifiers (and thus blank nodes) based on the alternation of universal and existential quantifiers:

(1) “Someone knows each professor” can be represented in RDFLog as

$$\exists stu \forall prof ((prof, \text{rdf:type}, \text{uni:professor}) \rightarrow (stu, \text{uni:knows}, prof)) \quad (1)$$

We call such rules $\exists\forall$ rules. Some approaches such as [5] are limited to rules of this form.

(2) Imagine, that we would like to state that each lecture must be “practiced” by another course (such as a tutorial or practice lab) without knowing more about that course. This statement can not be expressed by $\exists\forall$ rules. In RDFLog it can be represented as

$$\forall lec \exists crs ((lec, \text{rdf:type}, \text{uni:lecture}) \rightarrow (crs, \text{uni:practices}, lec)) \quad (2)$$

Such rules are referred to as $\forall\exists$ rules. Recent proposals for rule extensions to SPARQL are limited to this form, if they consider blank nodes in rule heads at all. The reason is that in SPARQL `CONSTRUCT` patterns a fresh blank node is constructed for each binding of the universal variables (cf. Section 10.2.1 in [8]).

(3) To the best of our knowledge, RDFLog is the first RDF query language that supports the third kind of rules, where quantifiers are allowed to alternate freely: This allows to express statements such as, for each lecture there is a course that “practices” that lecture and is attended by all students attending the lecture. This is represented in RDFLog as

$$\forall lec \exists crs \forall stu ((lec, \text{rdf:type}, \text{uni:lecture}) \wedge (stu, \text{uni:attends}, lec) \rightarrow (crs, \text{uni:practices}, lec) \wedge (stu, \text{uni:attends}, crs)) \quad (3)$$

In addition to flexible support for existential information through full quantifier alternation, RDFLog captures the essentials of RDF through two further characteristics: First, RDFLog is a closed RDF query language, i.e., the answer to an RDFLog program is again an RDF graph. Second, RDFLog can express the logical core of the RDFS semantics (ρ df from [2]).

Contributions. The paper is organised along the following contributions:

1. A rule-based RDF query language combining *recursion and free quantifier alternation*, called RDFLog (Section 3) is introduced.
2. The closed semantics of RDFLog (with or without a core fragment of the RDFS semantics) is introduced in terms of RDF entailment. (Section 3.3).
3. We show how this semantics can be implemented by a *reduction to the evaluation of a standard logic program* without existential quantifiers. (Section 4).

4. The experimental evaluation of a basic prototype shows that the reduction to standard logic programming easily competes with existing specialized RDF query engines even when considering only the restricted fragment of RDFLog equivalent to SPARQL (Section 5).

2 Preliminaries

2.1 Syntax and Semantics of RDF

In this paper, we adopt the notions of RDF vocabulary, RDF graph, (simple) RDF interpretation, and RDF entailment from [10].

Definition 1 (RDF Graph [10]). An RDF vocabulary \mathcal{V} consists of two disjoint sets called URIs \mathcal{U} and literals \mathcal{L} . The blank nodes \mathcal{B} is a set disjoint from \mathcal{U} and \mathcal{L} . An RDF graph is a set of RDF triples where an RDF triple is an element of $(\mathcal{U} \cup \mathcal{B}) \times \mathcal{U} \times (\mathcal{U} \cup \mathcal{L} \cup \mathcal{B})$. If $t = (s, p, o)$ is an RDF triple then s is the subject, p is the predicate, and o is the object of t .

The set \mathcal{L} of literals consists of three subsets, *plain literals*, *typed literals* and *literals with language tags*. In this work we consider only plain literals (and thus drop \mathcal{IL} , the interpretation function for typed literals, see Section 1.3 in [10], in the following definitions).

Definition 2 (RDF Interpretation [10]). An interpretation \mathcal{I} of an RDF vocabulary $\mathcal{V} = (\mathcal{U}, \mathcal{L})$ is a tuple $(\mathcal{IR}, \mathcal{LV}, \mathcal{IP}, \mathcal{IEXT}, \mathcal{IS})$ where \mathcal{IR} is a non-empty set of resources such that $\mathcal{L} \subseteq \mathcal{LV} \subseteq \mathcal{IR}$, \mathcal{IP} is a set of properties and $\mathcal{IEXT} : \mathcal{IP} \rightarrow 2^{\mathcal{IR} \times \mathcal{IR}}$, and $\mathcal{IS} : \mathcal{U} \rightarrow \mathcal{IR} \cup \mathcal{IP}$ are mappings.

Note that as \mathcal{IR} and \mathcal{IP} are not necessarily disjoint a same URI can be used both as a resource and a property. RDF interpretations are used to assign a truth value to an RDF graph.

RDF assigns a special meaning to a predefined vocabulary, called RDFS vocabulary. For example it is required that $\mathcal{IEXT}(\mathcal{IP}(\text{rdfs} : \text{subPropertyOf}))$ is transitive and reflexive. The formulation of these constraints on RDF interpretation makes use of a notion of a *class*. We have omitted this notion in the definition above for simplicity. The logical core of RDFS has been identified in [2], denoted as *pdf*. An RDF interpretation \mathcal{I} is a *pdf interpretation* if \mathcal{I} satisfied the constraints specified in Definition 3 in [2].

Definition 3 (Interpretation of an RDF Graph [10]). Let \mathcal{I} be the RDF (*pdf*) interpretation $(\mathcal{IR}, \mathcal{LV}, \mathcal{IP}, \mathcal{IEXT}, \mathcal{IS})$ and $A : \mathcal{B} \rightarrow \mathcal{IR}$ a mapping. Then $[\mathcal{I} + A](e) = a$ if e is the literal a , $[\mathcal{I} + A](e) = \mathcal{IS}(e)$ if e is a URI, $[\mathcal{I} + A](e) = A(e)$ if e is a blank node, and $[\mathcal{I} + A](e) = \text{true}$ if $e = (s, p, o)$ is an RDF triple over \mathcal{V} , $\mathcal{I}(p) \in \mathcal{IP}$ and $(\mathcal{I}(s), \mathcal{I}(o)) \in \mathcal{IEXT}(\mathcal{I}(p))$. Finally $\mathcal{I}(g) = \text{true}$ if there is a mapping $A : \mathcal{B} \rightarrow \mathcal{IR}$ such that $[\mathcal{I} + A](t) = \text{true}$ for all RDF triples $t \in g$.

The semantics of RDF is completed by the notion of entailment: An RDF graph g *RDF-entails* (*pdf-entails*) an RDF graph h if for all RDF (*pdf*) interpretations \mathcal{I} , $\mathcal{I}(h) = \text{true}$ if $\mathcal{I}(g) = \text{true}$ [10].

2.2 Logic and Logic Programming

We use formulas, terms, structures (rather than first-order interpretation), Herbrand structures, satisfaction \models , models, entailment \models , logic and Datalog programs, and *immediate consequence operator* T_P of a program P as common in logic and logic programming. In addition, we also consider infinite formulas: if Φ is a countably infinite set of formulas then $\bigwedge(\Phi)$ is a formula and if $\bar{x} = x_1, x_2, \dots$ is a countably infinite sequence of variables and φ is a formula then $\exists\bar{x}(\varphi)$ is a formula. We write $\varphi(\bar{x})$ to indicate that the free variables of a formula φ are among $\bar{x} = x_1, \dots, x_n$. Note, that we in fact use only a very limited form of infinite formulas (infinite conjunctions with only existential quantifiers).

3 Syntax and Semantics of RDFLog

3.1 The RDFLog Data Model

To make results from databases and logic programming accessible for RDF querying, we show that the semantics of RDF can be defined in terms of standard logic. In particular we show that RDF graphs can be translated to formulas so that logical entailment coincides with RDF entailment.

For any RDF vocabulary $\mathcal{V} = (\mathcal{U}, \mathcal{L})$ we define the alphabet $\Sigma_{\mathcal{V}} = \mathcal{U} \cup \mathcal{L} \cup \{T\}$ where \mathcal{U} and \mathcal{L} are constant symbols and T is an arbitrary ternary relation symbol.

Definition 4 (Canonical Formula of an RDF Graph). *Let $\mathbf{g} = \{t_1, \dots, t_n\}$ be an RDF graph over \mathcal{V} . The canonical formula of \mathbf{g} is the formula $\varphi_{\mathbf{g}} := \exists\bar{x}(\psi_1(\bar{x}) \wedge \dots \wedge \psi_n(\bar{x}))$ over $\Sigma_{\mathcal{V}}$ and variables from \mathcal{B} where $\psi_i = T(\mathbf{s}, \mathbf{p}, \mathbf{o})$ if $t_i = (\mathbf{s}, \mathbf{p}, \mathbf{o})$ and \bar{x} is the set of blank nodes occurring in \mathbf{g} .*

In [2] a sound and complete deductive system for *pdf* has been presented. It is easy to see that this deductive system corresponds to a finite set of Datalog rules Φ^{pdf} .

Proposition 1. *Let \mathbf{g}, \mathbf{h} be RDF graphs and $\varphi_{\mathbf{g}}, \varphi_{\mathbf{h}}$ their canonical formulas. Then \mathbf{g} RDF-entails \mathbf{h} iff $\varphi_{\mathbf{g}} \models \varphi_{\mathbf{h}}$ and \mathbf{g} pdf-entails \mathbf{h} iff $\varphi_{\mathbf{g}} \wedge \Phi^{pdf} \models \varphi_{\mathbf{h}}$.¹*

3.2 RDFLog Syntax

Definition 5 (Syntax of RDFLog Programs). *Let $\mathcal{V} = (\mathcal{U}, \mathcal{L})$ be an RDF vocabulary and Var a set of variables. An RDFLog atom over \mathcal{V} is an atom $T(t_1, t_2, t_3)$ where $t_1, t_2 \in (\mathcal{U} \cup Var)$ and $t_3 \in (\mathcal{U} \cup \mathcal{L} \cup Var)$. An RDFLog rule over \mathcal{V} is a formula*

$$\forall\bar{x}_1\exists\bar{y}_1\dots\forall\bar{x}_n\exists\bar{y}_n(\text{body}(\bar{x}) \rightarrow \text{head}(\bar{x}, \bar{y}))$$

¹ For proofs of theorems, lemmas, and proposition see the appendix of the online version [11].

over Σ_V and Var where $\bar{x} = \bar{x}_1, \dots, \bar{x}_n$ and $\bar{y} = \bar{y}_1, \dots, \bar{y}_n$ are finite sequences from Var and $body(\bar{x})$ and $head(\bar{x}, \bar{y})$ are finite conjunctions of RDFLog atoms. In addition we require that RDFLog rules are range restricted: if $x \in Var(head)$ is universal or there is an existential $y \in Var(head)$ such that y is in the scope of x , then $x \in Var(body)$. An RDFLog program over V is a finite set of RDFLog rules over V .

Observe that any finite RDF graph $g = \{t_1, \dots, t_n\}$ with blank nodes \bar{x} can be encoded into the RDFLog rule $\exists \bar{x} (true \rightarrow t_1 \wedge \dots \wedge t_n)$ where $true$ denotes the empty conjunction. As it makes the notation simpler we always assume that the input RDF graph is encoded into a rule in the RDFLog program. As there is only one predicate symbol (T) in an RDFLog program it is usually omitted.

3.3 RDFLog Semantics

It is not generally agreed upon what the semantics of a rule based RDF query language should be if existential variables are allowed in the head. In contrast, it is agreed that the semantics of a logic program with only universally quantified variables is its minimal Herbrand model.

The following RDFLog program illustrates why it is problematic to define the semantics of an RDF query language directly in terms of models. Let the *canonical structure* A_g of an RDF graph g be the structure over the domain of URIs, literals and blank nodes where (t_1, t_2, t_3) is true in A_g iff (t_1, t_2, t_3) is an RDF triple in g . As (2) is a fact in P and (1) is a rule in P , any canonical structure of an RDF graph that is a model of P must contain the triple ('Logic', uni:located_in, _:b) for some blank node _:b. Since this triple contains a literal in the subject position, it is not an RDF triple. This illustrates that P has no model that is the canonical structures of an RDF graph. Even if literals in subject position are allowed (as in SPARQL), a similar argument can be made with blank nodes in predicate position.

$$P = \{ \forall sem \exists rm \forall stu ((stu, uni:attends, sem) \rightarrow (sem, uni:located_in, rm) \wedge (stu, uni:knows, rm)), \quad (1)$$

$$true \rightarrow (uni:julie, uni:attends, 'Logic') \wedge (uni:john, uni:attends, uni:RDF) \} \quad (2)$$

$$\llbracket P \rrbracket \ni \{ (_ :b3, uni:located_in, _ :b1), (uni:julie, uni:knows, _ :b1), (uni:RDF, uni:located_in, _ :b2), (uni:john, uni:knows, _ :b2), (uni:julie, uni:attends, 'Logic'), (uni:julie, uni:attends, _ :b3), (uni:john, uni:attends, uni:RDF) \}$$

We deal with this problem by defining the semantics of RDFLog in terms of RDF entailment. More precisely we define the semantics of an RDFLog program P to be the set of all RDF graphs g that entail exactly the same RDF graphs as P (and satisfying in particular $P \models g$).

Definition 6 (Denotational Semantics of RDFLog). *Let P be an RDFLog program and RDF the set of RDF graphs. The denotational semantics $\llbracket P \rrbracket$ of P*

is the set $\llbracket P \rrbracket := \{g \in \text{RDF} \mid \forall h \in \text{RDF} (P \models \varphi_h \text{ iff } \varphi_g \models \varphi_h)\}$ where φ_g and φ_h are the canonical formulas of g and h respectively.

Observe that the semantics of an RDFLog program is an infinite set of possibly infinite RDF graphs. As we formalised RDF graphs as formulas, we have to consider the special kind of infinite formulas defined in section 2.2. Nonetheless it is immediate from the definition that the RDF graphs in $\llbracket P \rrbracket$ form an equivalence class under RDF entailment. Therefore any element of $\llbracket P \rrbracket$ characterizes the infinite set $\llbracket P \rrbracket$. In the next section we show how such a representative can be computed.

Observe that Φ^{pdf} encoded in RDFLog. Therefore it is up to the programmer to enclose Φ^{pdf} into P if the semantics of P is supposed to be aware of the *pdf* vocabulary.

4 Evaluation

The goal of this section is to show how the evaluation of an RDFLog program P can be done by first translating P into a logic program $s(P)$, using the well-studied notion of Skolemisation, and then evaluate this program $s(P)$ using standard technology. Two post processing steps (Unskolemisation and RDF normalization) make sure that the result is an RDF graph in the denotational semantics of P . After defining precisely each of the key steps of the operational semantics in Section 4.1, we show in Section 4.2 that the operational semantics achieves its goal as it is consistent with the denotational semantics of RDFLog.

4.1 Operational Semantics of RDFLog

Definition 7 (Skolemisation). Let Σ and Γ be disjoint alphabets, $\varphi = \forall \bar{x} \exists y (\psi)$ a formula over $\Sigma \cup \Gamma$ and $f \in \Gamma$. A Γ -Skolemisation step s_f maps φ to $s_f(\varphi) := \forall \bar{x} \psi \{y \leftarrow f(\bar{x})\}$. A Γ -Skolemisation s is a composition $s_{f_1} \circ \dots \circ s_{f_n}$ of Γ -Skolemisation steps such that f_i does not occur in $s_{f_{i+1}} \circ \dots \circ s_{f_n}(\varphi)$ and $s(\varphi)$ contains no existential variables. The definition of a Skolemisation is extended to sets in the usual way.

The Skolemisation of an RDFLog program P is equivalent to a range restricted logic program, which we denote by $s(P)$. Any logic programming engine can compute the minimal Herbrand model $M_{s(P)}$ of $s(P)$. The following logic program is the Skolemisation $s(P)$ of the RDFLog program P from Section 3.3 where s replaces the existential variable rm in P by the term $s_{rm}(sem)$.

$$\begin{aligned}
s(P) = & \{ \forall sem \forall stu ((stu, uni:attends, sem) \\
& \rightarrow (sem, uni:located_in, s_{rm}(sem)) \wedge (stu, uni:knows, s_{rm}(sem))), \\
& true \rightarrow (uni:julie, uni:attends, 'Logic') \wedge uni:john, uni:attends, uni:RDF) \} \\
\varphi_{M_{s(P)}} = & ('Logic', uni:located_in, s_{rm}('Logic')) \wedge (uni:julie, uni:knows, s_{rm}('Logic')) \\
& \wedge (uni:RDF, uni:located_in, s_{rm}(uni:RDF)) \wedge (uni:john, uni:knows, s_{rm}(uni:RDF)) \\
& \wedge (uni:julie, uni:attends, 'Logic') \wedge (uni:john, uni:attends, uni:RDF)
\end{aligned}$$

We define $\varphi_{M_S(P)}$ to be the conjunction of all ground atoms that are true in $M_S(P)$. However, $\varphi_{M_S(P)}$ might not be the canonical formula of an element of $\llbracket P \rrbracket$ for two reasons. First, the example shows that $\varphi_{M_S(P)}$ might contain atoms with skolem terms, such as $(\text{uni:RDF}, \text{uni:located_in}, s_{rm}(\text{uni:RDF}))$, which are not entailed by P . Second, $\varphi_{M_S(P)}$ can contain atoms that contain literals in subject or predicate position and blank nodes in predicate position. In the example the atom $(\text{'Logic'}, \text{uni:located_in}, s_{rm}(\text{'Logic'}))$ contains the literal 'Logic' in subject position.

We can avoid the first problem by “undoing” the Skolemisation: replacing each Skolem term in $\varphi_{M_S(P)}$ by a fresh, distinct blank node. We formalise this operation as the inverse of a Skolemisation called *Unskolemisation*.

Definition 8 (Unskolemisation). *Let Σ and Γ be disjoint alphabets and φ a ground, possibly infinite, and quantifier free formula over $\Sigma \cup \Gamma$. Let \bar{t} be the sequence of all ground terms $f(\bar{u})$ where f is in Γ and \bar{u} is a sequence of terms over $\Sigma \cup \Gamma$. Then the Γ -Unskolemisation u maps φ to $u(\varphi) := \exists \bar{x} (\varphi\{\bar{t} \leftarrow \bar{x}\})$, where \bar{x} is a sequence of fresh variables.*

To address the second issue, we remove all triples with literals or blank nodes in predicate position (no RDF graph may contain such a triple or any triple entailed by it). In addition we remove each triple t that contains a literal l in object position and add two triples t_1 and t_2 where t_1 is obtained from t by replacing an occurrence of a literal l in subject position by a fresh blank node b_l and t_2 is obtained from t by replacing all occurrences of l by b_l .

This is necessary to preserve information about the identity of domain elements that are denoted by blank nodes. For example observe that the RDF graph $\{(\text{uni:julie}, \text{uni:attends}, _:\text{b}), (_:\text{b}, \text{uni:located_in}, s_{rm}(\text{'Logic'}))\}$ follows from the RDFLog program P in Section 3.3. To maintain this information we need to insert the triple $(\text{uni:julie}, \text{uni:attends}, _:\text{b3})$ into $\llbracket P \rrbracket$. We formalise this step by defining the normalisation operator.

Definition 9 (Normalisation Operator). *Let φ be a formula of the form $\exists \bar{x} (a_1(\bar{x}) \wedge \dots \wedge a_n(\bar{x}))$ where each $a_i(\bar{x}) = T(t_1, t_2, t_3)$ for some $t_1, t_2, t_3 \in (\mathbf{U} \cup \mathbf{B} \cup \mathbf{L})$. Let $\mathbf{L}' \subseteq \mathbf{L}$ be the set of literals that occur in the first argument of an atom in φ . We define $\mu : \mathbf{U} \cup \mathbf{B} \cup \mathbf{L} \rightarrow \mathbf{U} \cup \mathbf{B} \cup \mathbf{L}$ to be the injection such that $\mu(t) = b$ for some fresh blank node b (not in φ) if $t \in \mathbf{L}'$ and $\mu(t) = t$ otherwise. Then $\Pi(\varphi) = \{\Pi(a_1(\bar{x})), \dots, \Pi(a_n(\bar{x}))\}$ and*

$$\Pi(T(t_1, t_2, t_3)) = \begin{cases} \top & \text{if } t_2 \in \mathbf{B} \cup \mathbf{L} \\ (\mu(t_1), t_2, t_3) \wedge (\mu(t_1), t_2, \mu(t_3)) & \text{otherwise} \end{cases}$$

The normalisation operator ensures that, though intermediary triples may contain blank nodes in predicate position (see [12] for examples where this is useful), the final answer of an RDFLog program never contains such triples. Armed with these notions of Skolemisation, Unskolemisation and Normalisation, we finally define the operational semantics of RDFLog as follows.

Definition 10 (Operational Semantics of RDFLog). *Let P be an RDFLog program over Σ , s a Γ -Skolemisation for P , and u an Γ -Unskolemisation. Then the operational semantics of P is $[P] := \Pi(u(\varphi_{M_s(P)}))$ where $\varphi_{M_s(P)}$ is as defined above: the conjunction of all ground atoms that are true in the minimal Herbrand model of $s(P)$.*

4.2 Properties of the Operational Semantics

Even though we do not require that elements of the denotational semantics $\llbracket P \rrbracket$ of an RDFLog program P are models of P it holds that $u(\varphi_{M_s(P)})$ has a canonical structure that is not only a model of P but even a universal model [9]. Thus if we allow literals in subject position and blank nodes in subject or predicate position, we can omit Π from the operational semantics and compute a model of P .

To formulate this more precisely, we define an *extended Herbrand structure* A over alphabet Σ and variables Var as a structure (D, Rel, Fun) where D is the set of (possibly non-ground) terms over Σ and Var , and every function f^A is defined by $f^A(t_1, \dots, t_n) = f(t_1, \dots, t_n)$. We extend the definition of Unskolemisation from formulas to extended Herbrand structures: if u is an Unskolemisation that replaces \bar{t} by \bar{x} then $u(M)$ is the extended Herbrand structure obtained from M by renaming the domain elements \bar{t} by \bar{x} .

Lemma 1. *Let P be an RDFLog program, $A_P = u(M_s(P))$ and $\varphi_P = u(\varphi_{M_s(P)})$. Then $A_P \models P$ and $P \models \varphi_P$.*

Intuitively, $A_P \models P$ means that φ_P captures all the information in P and $P \models \varphi_P$ means that it does not assert anything that is not asserted by P . From these two key observations, we can prove that the operational semantics of RDFLog is both sound and complete with respect to the denotational semantics.


Theorem 1. *Let P be an RDFLog program. Then $[P] \in \llbracket P \rrbracket$.*

5 Experimental Evaluation

The reduction of RDFLog to standard logic programs (Section 4) allows for a direct implementation of RDFLog on top of any logic programming or database engine that supports value invention and recursion. In the following, we compare experimentally the performance of a very simple prototype based on that principle with two of the more common SPARQL implementations. Our implementation of RDFLog uses a combination of Perl pre- and post-filters for Skolemisation, Unskolemisation, and normalisation of RDFLog programs and XSB Prolog to evaluate the Skolemised programs.

We compare our implementation with the ARQ SPARQL processor of Jena (Version 2.1) and the SPARQL engine provided by the Sesame RDF Framework. For Sesame, we choose the main-memory store as it is “by far the fastest type of repository that can be used” according to Sesame’s authors. With this store,

Fig. 1 Performance comparison on rule 1 (left) and on rule 2 (right)



Sesame becomes a main-memory, ad-hoc query engine just like RDFLog and ARQ. As common for ad-hoc queries we measure overall execution time including both loading of the RDF data and execution of the SPARQL or RDFLog query.

In the experiments we evaluate three different queries against an RDF graph consisting of Wikipedia data. The experiments have been carried out on a Intel Pentium M Dual-Core with 1.86 GHz, 1 MB cache and 2 GB main memory. For each setting, the running time is averaged over 25 runs. We compare the following rules:

- Rule 1: $\forall x \forall y ((x, \text{wiki:internalLink}, y) \rightarrow (x, \text{test:connected}, y))$
- Rule 2: $\forall x \forall y \exists z ((x, \text{wiki:internalLink}, y) \rightarrow (x, \text{test:connected}, z))$

Figure 1 compares the performance of RDFLog with that of ARQ and Sesame for rule 1 and rule 2 (we omit rule 3 as it is not expressible in SPARQL). Despite its light-weight, ad-hoc implementation, RDFLog outperforms ARQ and Sesame in this setting. The figures show moreover that also for ARQ and Sesame, blank node construction does not bear any significant additional computational effort.

6 Conclusion

Blank nodes are one of RDF's distinguishing features. Yet they have been entirely neglected or treated only in a limited fashion in previous approaches to RDF querying. With RDFLog we propose a simple, yet comprehensive extension of Datalog that covers all aspects of blank node *construction* that arise when combining RDF with rules. We show that such an extension, including the restrictions of RDF wrt. blank node occurrence can be treated in a semantics based purely on entailment. Furthermore, RDFLog easily incorporates (the logical core of) RDFS. This allows us to view RDFLog as a convenient vessel for classifying and comparing RDF query languages, similar to the role of Datalog for relational databases. In particular, we identify four classes of blank node support in a rule based RDF query language: no support, in the scope of *no* universal variable,

in the scope of *all* universal variables, or arbitrarily alternating with universal variables. Existing approaches fall in one of the three first classes, with RDFLog the first instance of the fourth class.

Though RDFLog is primarily designed as a logical foundation for RDF query languages, we also show that it is easily implemented on top of existing logic programming technology and that such an approach actually compares very well with existing SPARQL engines.

Acknowledgements. We would like to thank Michael Benedikt for the fruitful discussions on RDFLog and the help with this article.

References

1. Furche, T., Linse, B., Bry, F., Plexousakis, D., Gottlob, G.: RDF Querying: Language Constructs and Evaluation Methods Compared. In: Tutorial Lectures Int'l. Summer School 'Reasoning Web'. Volume 4126 of Lecture Notes in Computer Science., Springer Verlag (2006) 1–52
2. Muñoz, S., Pérez, J., Gutierrez, C.: Minimal Deductive Systems for RDF. In: Proc. European Semantic Web Conf. (ESWC). Volume 4519 of Lecture Notes in Computer Science., Springer Verlag (2007) 53–67
3. Polleres, A.: From SPARQL to Rules (and Back). In: Proc. Int'l. World Wide Web Conf. (WWW), New York, NY, USA, ACM (2007) 787–796
4. Sintek, M., Decker, S.: Triple—a Query, Inference, and Transformation Language for the Semantic Web. In: Proc. Int'l. Semantic Web Conf. (ISWC). (2002)
5. Yang, G., Kifer, M.: Reasoning about Anonymous Resources and Meta Statements on the Semantic Web. *Journal of Data Semantics* **1** (2003) 69–97
6. Schenk, S., Staab, S.: Networked Graphs: a Declarative Mechanism for SPARQL Rules, SPARQL Views and RDF Data Integration on the Web. In: Proc. Int'l. World Wide Web Conf. (WWW), New York, NY, USA, ACM (2008) 585–594
7. Gutierrez, C., Hurtado, C., Mendelzon, A.O.: Foundations of Semantic Web Databases. In: Proc. ACM Symp. on Principles of Database Systems (PODS), New York, NY, USA, ACM Press (2004) 95–106
8. Prud'hommeaux, E., Seaborne, A.: SPARQL Query Language for RDF. Proposed recommendation, W3C (2007)
9. Fagin, R., Kolaitis, P.G., Miller, R.J., Popa, L.: Data Exchange: Semantics and Query Answering. (2003) 207–224
10. Hayes, P., McBride, B.: RDF Semantics. Recommendation, W3C (2004)
11. Bry, F., Furche, T., Ley, C., Linse, B., Marnette, B.: RDFLog: It's like Datalog for RDF. Technical Report PMS-FB-2008-01, University of Munich (2008) <http://rdflog.com/publications/bry-rdflog-full.pdf>.
12. ter Horst, H.J.: Completeness, Decidability and Complexity of Entailment for RDF Schema and a Semantic Extension Involving the OWL Vocabulary. *Web Semantics: Science, Services and Agents on the World Wide Web* **3** (2005)