

Semantic Navigation on the Web of Data: Specification of Routes, Web Fragments and Actions

Valeria Fionda
KRDB, Free University of
Bozen-Bolzano, Italy

Claudio Gutierrez
DCC, Universidad de Chile,
Santiago, Chile

Giuseppe Pirró
KRDB, Free University of
Bozen-Bolzano, Italy

ABSTRACT

The massive semantic data sources linked in the Web of Data give new meaning to old features like navigation; introduce new challenges like semantic specification of Web fragments; and make it possible to specify actions relying on semantic data. In this paper we introduce a declarative language to face these challenges. Based on navigational features, it is designed to specify fragments of the Web of Data and actions to be performed based on these data. We implement it in a centralized fashion, and show its power and performance. Finally, we explore the same ideas in a distributed setting, showing their feasibility, potentialities and challenges.

Keywords

Navigation, Web of Data, Linked Data, Semantic Web

1. INTRODUCTION

Classically the Web has been modelled as a huge graph of links between pages [4]. This model included Web features such as links without labels and *only* generated by the owner of the page.¹ Although Web pages are created and kept distributively, their small size and lack of structure stimulated the idea to view searching and querying through single and centralized repositories (built from pages via crawlers). With the advent of the Web of Data, that is, semantic data at massive scale [3, 16], these assumptions, in general, do not hold anymore. First, links are semantically labelled (thanks to RDF triples) thus can be used to orient and control the navigation, are generated distributively and can be part of any data source. Hence, it has become a reality –using the words of Tim Berners-Lee– that *anyone* can say *anything* about *anything* and publish it *anywhere*. Second, data sources have a truly distributed nature due to their huge size, autonomous generation, and standard RDF structure. This makes inconvenient and impractical to reorganize them in central repositories as for Web pages.

¹Even though the spec. XLink [7] allows to define links in a third page, it was never used massively.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...\$10.00.

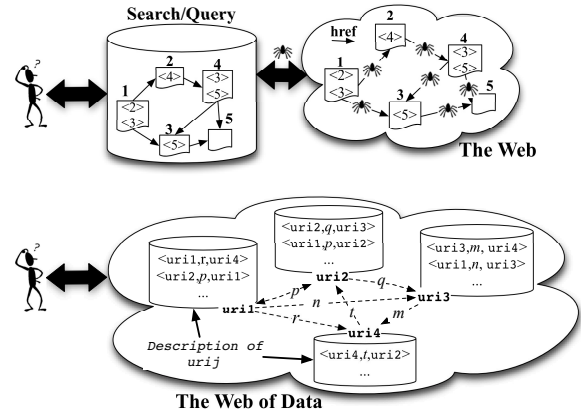


Figure 1: Classical Web versus Web of Data. Size, distributive character, and semantic description of data gives navigation a prominent role.

In this setting, navigation along the nodes of the Web of Data, using the semantics stored in each data source, becomes significant. To model these issues, rather than as a graph, the Web of Data is better represented as a set of nodes plus data describing their semantic structure “hanging” from each node (see Fig. 1). This model permits to better express the distributed creation and maintenance of data, and the fact that its structure is provided by dynamical and distributed data sources. In particular, it reflects the fact that at each moment of time, and for each particular agent, the whole network of data on the Web is unknown [19].

This new scenario calls for new models and languages to query and explore this semantic data space. In particular we highlight three functionalities: (1) a new type of navigation emerges as an important feature, in order to traverse sites and data sources; (2) closely tied to it, navigation charts or specifications, that is, semantic descriptions of fragments of the Web; (3) specification of actions one would like to perform over this data (*e.g.*, retrieving data, sending messages, etc.) also becomes relevant. Navigation, specifications of regions, and actions appear as part of the basic functionalities for exploring and doing data management over the Web of Data. Ideally, one would like to have a simple declarative language that integrates all of them.

In this paper we present such a language, which we call NAUTLOD, and show that it can be readily implemented on the current Linked Open Data (LOD) network [16]. In fact, we introduce the *swget* tool that exploits current Web protocols and work on LOD data. Finally, we explore its distributed version and implement an application as proof-of-concept to show its feasibility, potentialities and challenges.

information that these sites provide. Its basic navigational features are inspired both by `wget` and `XPath`, enhanced with semantic specifications, using SPARQL to filter paths, and with actions to be performed while navigating. We present a simple syntax, a formal semantics and a basic cost analysis.

(2) Second: *we implement a version of the language*, by developing the application `swget` that evaluates NAUTiLOD expressions in a centralized form (at the distinguished initial node). Being based on NAUTiLOD, `swget` permits to perform semantically-driven navigation of the Web of Data as well as retrieval actions. This tool relies on the computational resources of the initial node issuing the command and exploits the Web protocol HTTP. It is readily available on the current Linked Open Data (LOD) network. Its limitation is, of course, the scalability: the traffic of data involved could be high, making the navigation costly.

(3) Third: *we implement swget in a distributed environment*. Based on simple assumptions on third parties (a small application that each server should run to join it, and that in many ways extends the idea of current *endpoints*), we show the feasibility of such an application that simulates a travelling agent, and hint at the powerful uses it can have. From this proof-of-concept, we explore the potentialities of this idea and its challenges.

The paper is organized as follows. Section 2 provides a quick overview of the Web of Data. In Section 3 the NAUTiLOD language is introduced: syntax, semantics and its evaluation costs. In Section 4 `swget`, a centralized implementation of NAUTiLOD is introduced: its architecture, pseudo-code and experimental evaluation. Section 5 deals with the distributed version of `swget`, showing the feasibility and potentialities of this application. Section 6 discusses related work. Finally, in Section 7 we draw conclusions and delineate future work.

2. PRELIMINARIES: THE WEB OF DATA

This section provides some background on RDF and Linked Open Data (LOD) that are at the basis of the Web of Data. For further details the reader can refer to [3, 11].

RDF. The Resource Description Framework (RDF) is a metadata model introduced by the W3C for representing information about resources in the Semantic Web. RDF is built upon the notion of *statement*. A statement defines the *property* p holding between two resources, the *subject* s and the *object* o . It is denoted by $\langle s, p, o \rangle$, and thus called *triple* in RDF. A collection of RDF triples is referred to as an *RDF graph*. RDF exploits Uniform Resource Identifiers (URIs) to identify resources. URIs represent global identifiers in the Web and enable to access the *descriptions* of resources according to specific protocols (e.g., HTTP).

2.1 Web of Data - the LOD initiative

The LOD initiative leverages RDF to publish and interlinking resources on the Web. This enables a new (semantic) space called *Web of Data*. Objects in this space are *linked* and looked-up by exploiting (Semantic) Web languages and technologies. LOD is based on some principles, which can be seen more as best practices than formal constraints:

(1) Real world objects or abstract concepts must be assigned names on the form of URIs.

(2) In particular, HTTP URIs have to be used so that people can look them up by using existing technologies.

(3) When someone looks up a URI, associated information has to be provided in a standard form (e.g., RDF).

(4) Interconnections among URIs have to be provided by including references to other URIs.

An important notion in this context is that of dereferenceable URI. A dereferenceable URI, represents an identifier of a real world entity that can be used to retrieve a representation, by an HTTP GET, of the resource it identifies. The client can negotiate the format (e.g., RDF, N3) in which it prefers to receive the description.

2.2 Data in the LOD

Data in the LOD are provided by sites (i.e., servers), which cover a variety of domains. For instance, `dbpedia.org` or `freebase.org` provide cross-domain information, `geonames.org` publishes geographic information, `pubmed.org` information in the domain of life-science whereas `acm.org` covers information about scientific publications.

Theoretically in each server resides an RDF triple-store (or a repository of RDF data). In order to obtain information about the resource identified by a URI u , a client has to perform an HTTP GET u request. This request is handled by the Linked Data server, which answers with a set triples. This is usually said to be the *dereferencing* of u .

In the Web of Data, resources are not isolated from one another, in spirit with the fourth principle of LOD, but are linked. The interlinking of these resources and thus of the corresponding sites in which they reside forms the so called *Linked Open Data Cloud*².

3. A NAVIGATION LANGUAGE FOR THE WEB OF DATA

As we argued in the Introduction, there are data management challenges emerging in the Web of Data that need to be addressed. Particularly important are: (i) the specification of parts of this Web, thus of semantic fragments of it; (ii) the possibility to declaratively specify the navigation and exploit the semantics of data placed at each node of the Web; (iii) performing actions while navigating. To cope with this needs, this section presents a *navigation language* for the Web of Data, inspired by two non-related languages: `wget`, a language to automatically navigate and retrieve Web pages; and `XPath`, a language to specify parts of documents in the world of semi-structured data. We call it *Navigational language for Linked Open Data*, NAUTiLOD.

NAUTiLOD is built upon *navigational expressions*, based on regular expressions, filtered by tests using ASK-SPARQL queries (over the data residing in the nodes that are being navigated), and incorporating actions to be triggered while the navigation proceeds. NAUTiLOD allows to: (i) semantically specify collections of URIs; (ii) perform recursive navigation of the Web of Data, controlled using the semantics of the RDF data *hanging* from the URIs that are visited (that can be obtained by dereferencing these URIs); (iii) perform actions on specific URIs, as for instance, selectively retrieve data from them.

Before presenting the language, we present in Section 3.1 an abstract data model of the Web of Data. Then we present the syntax of NAUTiLOD (Section 3.2), and the formal semantics (Section 3.3). Finally, we provide a basic cost model for the complexity of evaluating NAUTiLOD expressions.

²<http://richard.cyganiak.de/2007/10/1od/>

3.1 Data model

We define a minimal abstract model of the Web of Data to highlight the main features required in our discussion.

Let \mathcal{U} be the set of all URIs and \mathcal{L} the set of all literals. We distinguish between two types of triples. *RDF links* $\langle s, p, o \rangle \in \mathcal{U} \times \mathcal{U} \times \mathcal{U}$ that encode connections among resources in the Web of Data. *Literal triples*, $\langle s, p, o \rangle \in \mathcal{U} \times \mathcal{U} \times \mathcal{L}$, which are used to state properties or features of the resource identified by the subject s . Note that the object of a triple, in the general case, can be also a blank node. However, here we will not consider them to simplify the presentation of the main ideas (note also that the usage of blank nodes is discouraged [16]). Let \mathcal{T} be the set of all triples in the Web of Data. The following three notions will be fundamental.

Definition 3.1 (WEB OF DATA \mathcal{T}). *Let \mathcal{U} and \mathcal{L} be infinite sets. The Web of Data (over \mathcal{U} and \mathcal{L}) is the set of triples $\langle s, p, o \rangle$ in $\mathcal{U} \times \mathcal{U} \times (\mathcal{U} \cup \mathcal{L})$. We will denote it by \mathcal{T} .*

Definition 3.2 (DESCRIPTION FUNCTION \mathcal{D}). *A function $\mathcal{D} : \mathcal{U} \rightarrow P(\mathcal{T})$ associates to each URI $u \in \mathcal{U}$ a subset of triples of \mathcal{T} , denoted by $\mathcal{D}(u)$, which is the set of triples obtained by dereferencing u .*

Definition 3.3 (WEB OF DATA INSTANCE \mathcal{W}). *A Web of Data instance is a pair $\mathcal{W} = \langle \mathcal{U}, \mathcal{D} \rangle$, where \mathcal{U} is the set of all URIs and \mathcal{D} is a description function.*

Note that not all the URIs in \mathcal{U} are dereferenciable. If a URI $u \in \mathcal{U}$ is not dereferenciable then $\mathcal{D}(u) = \emptyset$.

3.2 Syntax

NAUTILOD provides a mechanism to declaratively: (i) define *navigational expressions*; (ii) allow *semantic control* over the navigation via test queries; (iii) *retrieve data* by performing actions as side-effects along the navigational path.

The navigational core of the language is based on regular path expressions, pretty much like Web query languages and XPath. The semantic control is done via existential tests using ASK-SPARQL queries. This mechanism allows to redirect the navigation based on the information present at each node of the navigation path. Finally, the language allows to command actions during the navigation according to decisions based on the original specification and the local information found.

<code>path ::=</code>	<code>pred (pred)⁻¹ action path/path (path)? (path)* (path path) path[test]</code>
<code>pred ::=</code>	<code><RDF predicate> <-></code>
<code>test ::=</code>	<code>ASK-SPARQL query</code>
<code>action ::=</code>	<code>procedure[Select-SPARQL query]</code>

Table 1: Syntax of the NautiLOD language.

The syntax of the language NAUTILOD is defined according to the grammar reported in Table 1. The language is based on *Paths Expressions*, that is, concatenation of base-case expressions built over *predicates*, *tests* and *actions*. The language accepts concatenations of basic and complex types of expressions. Basic expressions are predicates and actions; complex expressions are disjunctions of expressions; expressions involving a number of repetitions using the features of regular languages; and expressions followed by a test. The building blocks of a NAUTILOD expression are:

1. *Predicates*. The base case. `pred` can be an RDF predicate or the wildcard `<->` used to denote *any* predicate.
2. *Test Expressions*. A `test` denotes a query expression. Its base case is an ASK-SPARQL query.
3. *Action Expressions*. An `action` is a procedural specification of a command (e.g., send a notification message, PUT and GET commands on the Web, etc.), which obtains its *parameters* from the data source reached during the navigation. It is a side-effect, that is, it does not influence the subsequent navigation process.

If restricted to (1) and (2), NAUTILOD can be seen as a declarative language to describe portions of the Web of Data, i.e., set of URIs conform to some semantic specification.

3.3 Semantics

NAUTILOD expressions are evaluated against a Web of Data instance \mathcal{W} and a URI u indicating the starting point of the evaluation. The meaning of a NAUTILOD expression is a set of URIs defined by the expression plus a set of actions produced by the evaluation of the expression. The resulting set of URIs are the leaves in the paths according to the NAUTILOD expression, originating from the seed URI u .

For instance, the expression `type`, evaluated over u , will return the set of URIs u_k reachable from u by “navigating” the predicate `type`, that is, by inspecting triples of the form $\langle u, \text{type}, u_k \rangle$ included in $\mathcal{D}(u)$. Similarly, the expression `type[q]` will filter, from the results of the evaluation of `type`, those URIs u_k for which the query q evaluated on their descriptions $\mathcal{D}(u_k)$ is true. Finally, the evaluation of an expression `type[q]/a` will return the results of `type[q]` and perform the action `a` (possibly using some data from $\mathcal{D}(u_k)$).

The formal semantics of NAUTILOD is reported in Table 2. The fragment of the language without actions follows the lines of formalization of XPath by Wadler [27]. Actions are treated essentially as side-effects and evaluated while navigating. Given an expression, a Web of Data instance $\mathcal{W} = \langle \mathcal{U}, \mathcal{D} \rangle$, and a seed URI u the semantics has the following modules:

- $E[\text{path}](u, \mathcal{W})$: Evaluates the set of URIs selected by the navigational expression `path` starting from the URI u in the Web of Data instance \mathcal{W} . Additionally, it collects the actions associated to each of such URIs.
- $U[\text{path}](u, \mathcal{W})$: Defines the set of URIs specified by the expression `path` when forgetting the actions.
- $A[\text{path}](u, \mathcal{W})$: Executes the actions specified by the evaluation of the navigational expression `path`.
- $Sem[\text{path}](u, \mathcal{W})$: Outputs the meaning of the expression `path`, namely, the ordered pair of two sets: the set of URIs specified by the evaluation of `path`; and the set of actions performed according to this information.

Note on some decisions made: Any sensible real implementation can benefit from giving an order to the elements of the output action set. As far as the formal semantics, at this stage we assumed that actions are independent from one another and that the world \mathcal{W} is static during the evaluation (to avoid to overload our discussion with the relevant issue of synchronization, that is at this point orthogonal to the current proposal). Thus, we decided to denote the actions

$$\begin{aligned}
E[\langle p \rangle](u, \mathcal{W}) &= \{(u', \perp) \mid \langle u, \langle p \rangle, u' \rangle \in \mathcal{D}(u)\} \\
E[(\langle p \rangle)^{-1}](u, \mathcal{W}) &= \{(u', \perp) \mid \langle u', \langle p \rangle, u \rangle \in \mathcal{D}(u)\} \\
E[\langle _ \rangle](u, \mathcal{W}) &= \{(u', \perp) \mid \exists \langle p \rangle, \langle u, \langle p \rangle, u' \rangle \in \mathcal{D}(u)\} \\
E[\text{act}](u, \mathcal{W}) &= \{(u, \text{act})\} \\
E[\text{path}_1/\text{path}_2](u, \mathcal{W}) &= \{(u'', a) \in E[\text{path}_2](u', \mathcal{W}) : \exists b, (u', b) \in E[\text{path}_1](u, \mathcal{W})\} \\
E[(\text{path})?](u, \mathcal{W}) &= \{(u, \perp)\} \cup E[\text{path}](u, \mathcal{W}) \\
E[(\text{path})^*](u, \mathcal{W}) &= \{(u, \perp)\} \cup \bigcup_{i=1}^{\infty} E[\text{path}_i](u, \mathcal{W}) \mid \text{path}_1 = \text{path} \wedge \text{path}_i = \text{path}_{i-1}/\text{path} \\
E[\text{path}_1|\text{path}_2](u, \mathcal{W}) &= E[\text{path}_1](u, \mathcal{W}) \cup E[\text{path}_2](u, \mathcal{W}) \\
E[\text{path}[\text{test}]](u, \mathcal{W}) &= \{(u', a) \in E[\text{path}](u, \mathcal{W}) : \text{test}(u') = \text{true}\} \\
\\
U[\text{path}](u, \mathcal{W}) &= \{v : \exists a, (v, a) \in E[\text{path}](u, \mathcal{W})\} \\
A[\text{path}](u, \mathcal{W}) &= \{\text{Exec}(a, v) : (v, a) \in E[\text{path}](u, \mathcal{W})\} \\
\\
\text{Sem}[\text{path}](u, \mathcal{W}) &= (U[\text{path}](u, \mathcal{W}), A[\text{path}](u, \mathcal{W}))
\end{aligned}$$

Table 2: Semantics of NautiLOD. The semantics of an expression is composed of two sets: (1) the set of URIs of \mathcal{W} satisfying the specification; (2) the actions produced by the evaluation of the specification. $\text{Exec}(a, u)$ denotes the execution of action a over u . \perp indicates the empty action (i.e., no action).

produced by the evaluation of an expression as a set. It is not difficult to see that one could have chosen a list as the semantics for output actions.

3.4 Evaluation of Costs and Complexity

We present a general analysis of costs and complexity of the evaluation of NAUTiLOD expressions over a Web of Data instance \mathcal{W} . We can separate the costs in three parts, where E are expressions, \bar{E} action-and-test-free expressions, A actions and T tests:

$$\text{cost}(E, \mathcal{W}) = \text{cost}(\bar{E}, \mathcal{W}) + \text{cost}(A) + \text{cost}(T). \quad (1)$$

Since actions do not affect the navigation process we can treat their cost separately. Besides, in our language, tests are ASK-SPARQL queries having a different structure from the pure navigational path expressions of the language. Even in this case we can treat their cost independently.

Actions. NAUTiLOD is designed for acting on the Web of Data. In this scenario, the cost of actions has essentially two components: *execution* and *transmission*. The execution cost boils down to the cost of evaluating the SELECT SPARQL query that gives the action’s parameters. As for transmission costs, a typical example is the `wget` command, where the cost is the one given by the GET data command.

Action-and-test-free. This fragment of NAUTiLOD can be considered essentially as the PF fragment of *XPath* (location paths without conditions), that is well known to be (with respect to combined complexity) **NL**-complete under **L**-reductions (Thm. 4.3, [10]). The idea of the proof is simple: membership in **NL** follows from the fact that we can guess the path while we verify it in time **L**. The hardness essentially follows from a reduction from the directed graph reachability problem. Thus we have:

THEOREM 3.4. *With respect to combined complexity, the action-and-test-free fragment of NAUTiLOD is **NL**-complete under **L**-reductions.*

Combined refers to the fact that the input parameters are the expression size and the data size. Note that what really matters is not the whole Web (the data), but only the set of nodes reachable by the expression. Thus it is more precise to speak of expression size plus set-of-visited nodes size. The worst case is of course the whole size of the Web.

Tests. The evaluation of tests (i.e., ASK-SPARQL queries) has a cost. This cost is well known and one could choose particular fragments of SPARQL to control it [21]. However, tests will possibly reduce the size of the set of nodes visited during the evaluation. Thus the $\text{cost}(\bar{E}, \mathcal{W})$ has to be reduced to take into account the effective subset of nodes reachable thanks to the filtering performed by the tests. Let \mathcal{W}_T be \mathcal{W} when taking into account this filtering. We have:

$$\text{cost}(E, \mathcal{W}) = \text{cost}(\bar{E}, \mathcal{W}_T) + \text{cost}(A) + \text{cost}(T). \quad (2)$$

Section 4.2 will discuss some examples on real world data by underlining the contribution of each component of the cost.

Final Considerations. In a distributed setting, with partially unknown information and a network of almost unbound size, the notion “cost of evaluating an expression e ” appears less significant than in a controlled centralized environment. In this scenario, a more pertinent question seems to be: “given an amount of resources r and the expression e , how much can I get with r satisfying e ?”. This calls for optimizing (according to some parameters) the navigation starting from a given URI u , according to equation (2).

4. IMPLEMENTATION OF NAUTiLOD

This section deals with `swget`, a tool implementing NAUTiLOD. The tool `swget` implements all the navigational features of NAUTiLOD, a set of actions centred on retrieving data, and adds (for practical reasons) a set of ad-hoc options for further controlling the navigation from a network-oriented perspective (e.g., size of data transferred, latency time) that today’s are not yet found as RDF statements.

`swget` has been implemented in Java and is available as: (i) a developer release, which includes a command-line tool that is easily embeddable in custom applications; (ii) an end user release, which features a GUI. Further details, examples, the complete syntax along with the downloadable versions are available at the `swget`’s Web site³.

4.1 Architecture

The high level architecture of `swget` is reported in the left part of Fig. 3. The *Command interpreter* receives the input, i.e., a seed URI, a NAUTiLOD expression and a set of options.

³<http://swget.wordpress.com>

The input is then passed to the *Controller* module, which checks if a network request is admissible and possibly passes it to the *Network Manager*. A request is admissible if it complies with what specified by the NAUTLOD expression and with the network-related navigation parameters (see Section 4.1.1). The *Network Manager* performs HTTP GET requests to obtain streams of RDF data. These streams are processed for obtaining Jena RDF models, which will be passed to the *Link Extractor*. The *Link Extractor* takes in input an automaton constructed by the NAUTLOD *interpreter* and selects a subset of outgoing links in the current model according to the current state of the automaton. The set is given to the *Controller Module*, which starts over the cycle. The execution will end either when some navigational parameter is satisfied or when there are no more URIs to be dereferenced.




Figure 3: swget architecture and scenario.

4.1.1 Network-based controlled navigation

NAUTLOD is designed to semantically control the navigation. However, it can be the case that a user wants to control the navigation also in terms of network traffic generated. A typical example is a user running *swget* from a mobile device with limited Internet capabilities. This is why *swget* includes features to add more control to the navigation through the parameters reported in Table 3. Each option is given in input to *swget* as a pair $\langle \text{param}, \text{value} \rangle$.

Table 3: Network params to control the navigation

Parameter	Value	Meaning
maxDerTriples	int	max. number of triples allowed in each dereferencing
saveGraph	boolean	Save the graphs dereferenced
maxSize	int	traffic limit (in MBs)
timeoutDer	long	connection time-out
timeout	long	total time-out
domains	List<String>	trusted servers

To illustrate a possible scenario where the navigation can be controlled both from a semantic and network-based perspective consider, the following example.

Example 4.1. (Controlled navigation) Find information about Rome, starting from its definition in DBpedia and includes other possible definitions of Rome linked to DBpedia but only if their description contains less than 500 triples and belongs to DBpedia, Freebase or The New York Times.

```
swget < dbp:Rome> (<owl:sameAs>)* -saveGraph
      -domains {dbpedia.org,rdf.freebase.com,
               data.nytimes.com} -maxDerTriples 500
```

The command, besides the NAUTLOD expression, contains the `-domains` and `-maxDerTriples` parameters to control the navigation on the basis of the trust toward information providers and the number of triples, respectively.

4.2 Evaluating NAUTLOD expressions

Given a NAUTLOD expression e it is possible to build an automaton that can recognize NAUTLOD expressions. The transitions between states of the automaton implements the navigation process.

4.2.1 The swget Navigation Algorithm

The *swget* controlled navigation algorithm is reported in Algorithm 1. Moreover, Table 4 describes the high level primitives used in the pseudo-code to interact with the automaton.

Algorithm 1: swget pseudo-code

```
Input : e=NAUTLOD expression; seed=URI; par=Parms<n,v>
Output: set of URIs and literals conform to e and par;

1 a = buildAutomaton(e);
2 addLookupPair(seed, a.getInitial());
3 while (∃ p=<uri,state> to look up and checkNet(par)=OK) do
4   desc=getDescription(p.uri);
5   if (a.isFinal(p.state)) then
6     addToResult(p.uri);
7   if (not alreadyLookedUp(p)) then
8     setAlreadyLookedUp(p);
9     if (t=getTest(p.state)≠∅ and evalT(t,desc)=true) then
10      s=a.nextState(p.state,t);
11      addLookupPair(p.uri,s);
12     if (act=getAction(p.state)≠∅) then
13       if (evalA(act.test,desc)) then exeC(act.cmd);
14       s=a.nextState(p.state,act);
15       addLookupPair(p.uri,s);
16     out=navigate(p,a,desc);
17     for (each URI pair p'=<uri,state> in out) do
18       addLookupPair(p');
19     for (each literal pair lit=<literal,state> in out) do
20       if (a.isFinal(lit.state)) then
21         addToResult(lit.literal);
22 return Result;
```

Function navigate(exp,a,desc)

Output: List of <uri,state> and <literal,state>

```
1 for (each pred in a.nextP(p.state)) do
2   nextS=a.nextState(p.state,pred);
3   query= "SELECT ?x WHERE
4     {{ ?x pred p.uri} UNION{ p.uri pred ?x}}";
5   for (each res in evalQ(query, desc)) do
6     addOutput(res,nextS);
7 return Output;
```

Table 4: Primitives for accessing the automaton.

Primitive	Behaviour
getInitial()	returns the initial state q_0
nextP(q)	returns the set $\{\sigma \mid \delta(q, \sigma) = q_1\}$ of tokens (i.e., predicates) enabling a transition from q to q_1
getTest(q)	returns the test to perform into the current automaton state
getAction(q)	returns the action to perform into the current automaton state
nextState(q,σ)	returns the state that can be reached from q by the token σ
isFinal(q)	returns TRUE if q is an accepting state

The algorithm takes as input a *seed* URI, a NAUTLOD expression and a set of network parameters, and returns a set of URIs and literals conform to the expression and the network parameters. For each URI involved in the evaluation, possible tests (line 9) and actions (line 12) are considered.

The procedure *navigate* is exploited to extract links (line 3) from a resource identified by $p.uri$ toward other resources. According to the Linked Data initial proposal [2] [section on browsable graphs] $p.uri$ may appear either as the *subject* or the *object* of each triple.




Figure 4: Evaluation of swget. Each expression has been executed 4 times. Average results are reported.

4.3 Experimental Evaluation

To show real costs of evaluating the different components of swget expressions over real-world data, we choose two complex expressions (shown in Fig. 5) to be evaluated over the Linked Open Data network. We report the results of swget in terms of execution time (t), URIs dereferenced (d) and number of triples retrieved (n). Each expression has been divided in 5 parts (i.e., $\sigma_i, i \in \{1..5\}$). They have been executed as whole (i.e., $\sigma_{[1-5]}$) and as action-and-test-free expressions (i.e., $\sigma_{[noAT]}$), which correspond to \bar{E}_1 and \bar{E}_2 , respectively (see Section 3.4). Moreover, the various sub-expressions (i.e., $\sigma_{[1-i]}, i \in \{1..4\}$) have also been executed. This leads to a total of 12 expressions. For each expression, the corresponding sub-Web has been locally retrieved. That is, for each reachable URI the corresponding RDF graph has been locally stored. The aim of the evaluation is to investigate how the various components in the cost model presented in Section 3.4 affect the parameters t , d and n .

seed URI: = http://dbp:Stanley_Kubrick	E₁
σ_1 : <dbpo:influenced><3>	
σ_2 : [ASK ?p <dbpo:birthData> ?y.FILTER(?y > 1961-01-01)]	
σ_3 : {sendEmail(?p)[SELECT ?p WHERE {?x <foaf:name> ?p.}]}	
σ_4 : <dbpo:director>	
σ_5 : <owl:sameAs>?	
seed URI: = http://dbp:Italy	E₂
σ_1 : <dbpo:homeTown>	
σ_2 : [ASK ?person <rdf:type> <dbpo:Person>. ?person <rdf:type> <dbpo:MusicalArtist>]	
σ_3 : <dbpo:birthPlace>	
σ_4 : [ASK ?t <dbpo:populationTotal> ?p. FILTER(?p < 15000)]	
σ_5 : <owl:sameAs>*	

Figure 5: Expressions used in the evaluation

The results of the evaluation, in logarithmic scale, are reported in Fig. 4(a)-(c). In particular, in the X axis are reported from left to right: the 4 sub-expressions, the full expression (i.e., $\sigma_{[1-5]}$) and the action-and-test-free expression (i.e., $\sigma_{[noAT]}$). Note that in some cases, the number of results is higher than the number of dereferenced URIs reported because not all the results were dereferenceable URIs.

The first expression (E_1) starts by finding people influenced by Stanley Kubrick up to a level 3 (subexpr. $\sigma_{[1-1]}$). This operation requires about 61 secs., for a total of 221 URIs dereferenced. On the description of each of these 221 URIs, an ASK query is performed to select only those entities that were born after the 1961 (subexpr. $\sigma_{[1-2]}$). The execution time of the queries is of about 4 secs. (i.e., ≈ 0.02 secs., per query). Hence, 31 entities have been selected. At this point, an action is performed on the descriptions of

these 31 entities by selecting their <foaf:name> to be sent via email (subexpr. $\sigma_{[1-3]}$). In total, the select, the rendering of the results in an HTML format and the transmission of the emails cost about 25 secs. The navigation continues from the 31 entities before the action to get movies through the property <dbpo:director> (subexpr. $\sigma_{[1-4]}$). The cost is of about 34 secs., for a total of 136 movies. Finally, for each movie only one level of possible additional descriptions is searched by the <owl:sameAs> property (the whole expr. $\sigma_{[1-5]}$) whose cost is 1638 secs., for a total of 409 new URIs available from multiple servers (e.g., linkedmdb.org, freebase.org) of which only 289 were dereferenceable.

By referring to the cost model in Section 3.4 we have that $cost(E_1, \mathcal{W}) = cost(\bar{E}_1, \mathcal{W}_{T_1}) + cost(A_1) + cost(T_1) = 1763$. Here, the factor $cost(A_1) \approx 25$ secs., whereas $cost(T_1) \approx 4$ secs., and $cost(\bar{E}_1, \mathcal{W}_{T_1}) \approx 1738$ secs. If we consider the test-and-action-free expression executed over the whole Web of Data (i.e., \mathcal{W}), we have that $cost(\bar{E}_1, \mathcal{W}) \approx 20018$ secs. Note that the ASK queries costs about 4 secs., and permits to reduce the portion of the Web of Data navigated by \bar{E}_1 , which enables to save about $20018 - 1738 = 18280$ secs. Such a larger difference in the execution times is justified by the fact that the 222 initial URIs, selected by $\sigma_{[1-1]}$ are not filtered in the case of (\bar{E}_1, \mathcal{W}) and then cause a larger amount of paths to be followed at the second level. Indeed, the total number of dereferenced URIs for (\bar{E}_1, \mathcal{W}) is 6053 while for $(\bar{E}_1, \mathcal{W}_{T_1})$ is 646 with about 660K triples retrieved in the first case and 125K in the second case.

The second expression (E_2) starts by navigating the property <dbpo:homeTown> to find entities living in Italy (subexpr. $\sigma_{[1-1]}$) with an execution time of about 84 secs., and a total of 400 dereferenced URIs, one seed and 399 URIs of entities. On the description of each of these 399 URIs, an ASK query filters entities that are of type <dbpo:Person> and <dbpo:MusicalArtist> (subexpr. $\sigma_{[1-2]}$). Hence, 399 ASK queries are performed for a total of about 3.8 secs., with an average time per query of 0.01 secs., to select 156 entities. For these entities, the navigation continues through the property <dbpo:birthPlace> to find the places where these people were born (subexpr. $\sigma_{[1-3]}$), which costs about $101 - 87 = 14$ secs. In total, 43 new URIs have been reached. The navigation continues with a second ASK query to select only those places in which live less than 15000 habitants (subexpr. $\sigma_{[1-4]}$). The cost of performing 43 ASK queries on the results of the previous step is of about 3 secs. Here 5 places are selected. Finally, for each of the 5 places additional descriptions are searched by navigating the <owl:sameAs> property (the whole expr. $\sigma_{[1-5]}$). This allows to reach a total of 29

URIs, some of which are external to `dbpedia.org`. The cost for this operation is of about 57 secs.

As for the cost, we have $cost(E_2, \mathcal{W}) = cost(\overline{E}_2, \mathcal{W}_{T_2}) + cost(A_2) + cost(T_2) = 161$. The factor $cost(A_2) = 0$ since E_2 does not contain any action whereas $cost(T_2) \simeq 6$ secs. Hence, $cost(\overline{E}_2, \mathcal{W}_{T_2}) = 155$ secs. The cost of the test-and-action-free expression (i.e., \overline{E}_2) over \mathcal{W} is $cost(\overline{E}_2, \mathcal{W}) \simeq 1600$ secs., for a total of 1277 dereferenced URIs. This is because the expression is not selective since it performs a sort of "semantic" crawling only based on RDF predicates. In fact, the number of triples retrieved (see Fig. 4(c)) is almost three times higher than in the case of the expression with tests. By including the tests, the evaluation of \overline{E}_2 is 1445 secs., faster.

5. A PROPOSAL FOR DISTRIBUTED `swget`

This section presents an overview of Distributed `swget` (`Dswget`), which has the peculiarity that the processing NAUTILOD expressions occurs in a cooperative manner among LOD information providers. The tool has been implemented and tested on a local area network.

5.1 `Dswget`: making LOD servers cooperate

`swget` enables controlled navigation but it heavily relies on the client that initiates the request. However, one may think of the Linked Data servers storing RDF triples as to peers in a Peer-to-Peer (P2P) network, where links are given by URIs in RDF triples. For instance `(dbp:Rome, owl:sameAs, fb:Rome)` links `dbpedia.org` with `freebase.org`. Indeed, there are some differences w.r.t. a traditional P2P network. First, Linked Data servers are less volatile than peers. Second, it is reasonable to assume that the computational power of Linked Data servers is higher than that of a traditional peers. This enables to handle a higher number of connections with the associated data.

Our proposal is to leverage the computational power of servers in the network to cooperatively evaluate `swget` commands. This enables to drastically reduce the amount of data transferred. In fact, data is not transferred from servers to the client that initiates the request (in response to HTTP GETs). Servers will exchange `swget` commands plus some metadata and operate on their data locally. This can be achieved by installing on each server in the network a `Dswget` engine and coordinating the cooperation by an ad-hoc distributed algorithm.

Table 5: Primitives of `Dswget`

Primitive	Behaviour
<code>sendResults</code>	sends to the original client (partial) results, which are URIs (line 14) and literals (line 18)
<code> fwdToServers</code>	forwards to other servers, the initial client address, the NAUTILOD expression and a set of pairs <code><URI, A State></code> . For each pair, the computation on a URI will be started from the corresponding <code>A state</code>

A `Dswget` command is issued by a `Dswget` client to the server to which the *seed* URI belongs. Each server involved in the computation will receive, handle and forward commands and results by using the Procedure `handle`. Note that in this procedure there are calls to some primitives reported in Table 4 and to the function `navigate` described in Section 4.2.1. The specific primitives needed by `Dswget` are reported in Table 5.

Procedure `handle(client_id, e, URIs, metadata)`

```

Input: client_id=address of the client; e=NAUTILOD
expression; URIs=set of pairs <URI, A state>;
metadata=additional data (e.g., current state of the
automaton, request id)

1 a=buildAutomaton(e);
2 for (each p=<uri, state> in URIs) do
3   desc=getDescription(p.uri); //local call no deref. needed
4   if (not alreadyLookedUp(p)) then
5     setAlreadyLookedUp(p);
6     if (t=getTest(p.state)≠∅ and evalT(t,desc)=true) then
7       s=a.nextState(p.state,t);
8       addLookupPair(p.uri,s);
9     if (act=getAction(p.state)≠∅) then
10      if (evalA(act.test,desc) then exeC(act.cmd);
11      out=navigate(p,a,desc);
12      for (each URI pair p'=<uri, state> in out) do
13        if (a.isFinal(p'.state) then
14          addToResults(p'.uri);
15          else addLookupPair(p');
16        for (each literal pair lit=<literal, state> in out) do
17          if (a.isFinal(lit.state) then
18            addToResult(lit.literal);
19      sendResults(client_id);
20      fwdToServers(client_id,e);

```

5.1.1 A Running example

To see an example of how `Dswget` works, consider the following request originated from a `Dswget` client:

Example 5.1. (`Dswget`) *Starting from DBPedia, find cities with less than 15000 persons, along with their aliases, in which musicians, currently living in Italy, were born.*

The `Dswget` command is reported in Fig. 6, which also reports a possible `Dswget` interaction scenario. On each linked data server a `Dswget` engine has been installed. Each server exposes a set of dereferenceable URIs for which the corresponding RDF descriptions are available. RDF data enables both internal references (e.g., `dbp:Rome` and `dbp:uri1`) and external ones (e.g., `fb:Enrico` and `geo:Paris`).

In Fig. 6 references between URIs are represented by dotted arrows. When not explicitly mentioned, it is assumed that the reference occurs on a generic predicate. The automaton associated to this expression, having `q4` and `q5` as accepting states, is also reported. The state(s) of the automaton on which a server is operating is(are) reported in grey. `Dswget` protocol messages have been numbered to emphasize the order in which they are exchanged.

The command along with some metadata (e.g., the address of the client) is issued by the client's `Dswget` engine toward the server to which the seed URI belongs (i.e., `dbpedia.org` in this example). The `Dswget` engine at this server, after locally building the automaton, starts the processing of the NAUTILOD expression at the state `q0`. It obtains from its local RDF store, the description of Rome `D(dbp:Rome)` and looks for URIs having `dbpo:hometown` as a predicate. In Fig 6, the URI `fb:Enrico` satisfies this pattern. The `Dswget` engine at `dbpedia.org` performs the first transition of state, that is, $\delta(q0, \sigma1) = q1$. The automaton does not reach a final state, and then the process has to continue. Since the URI `fb:Enrico` belongs to another server, the `Dswget` engine at `dbpedia.org`, communicates with that at `freebase.org` by seeding the initial NAUTILOD expression, the URI for which `freebase.org` is involved in the computation (i.e., `fb:Enrico`) and the current state of the automaton. In the case in which multiple URIs have to be sent, they are packed together in a unique message.

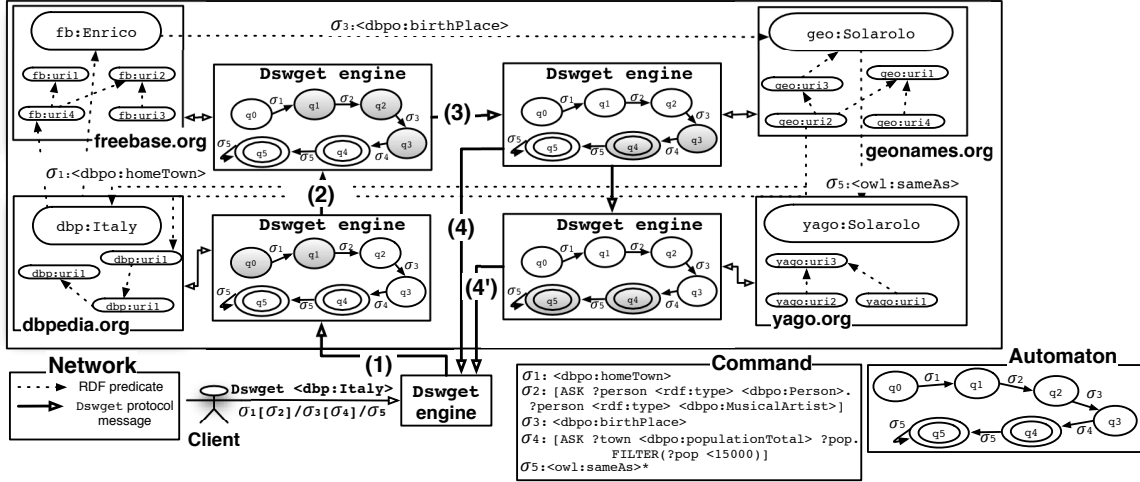


Figure 6: Distributed Dswget interaction scenario.

With a similar reasoning the request reaches the Dswget engine at `geonames.org`, which checks if it is possible to reach the next state of the automaton starting from the URI passed by `freebase.org`. It has to check on $\mathcal{D}(\text{geo:Solarolo})$ if the query represented by σ_4 can be satisfied, that is, whether this city has less than 15K inhabitants. Then, the state q_4 is reached, which is a final state. The Dswget engine at `geonames.org` contacts directly the Dswget engine of the client that issued the request and send the result (i.e., the URI `geo:Solarolo`). The address of the client is passed at each communication among Dswget engines.

Note that the automaton has another final state, that is, q_5 that can be reached if there exist some triples in $\mathcal{D}(\text{geo:Solarolo})$ having an `owl:sameAs` predicate. Such a triple is $(\text{geo:Solarolo}, \text{owl:sameAs}, \text{yago:Solarolo})$. Therefore, the Dswget engine at `geonames.org` sends to the engine at `yago.org` the URI in the object of this triple, the expression and the current state of the automaton. Here, as the automaton is in a final state, the Dswget engine sends to the client the result and continues the process. In this case since in $\mathcal{D}(\text{yago:Solarolo})$ there are no more triples having `owl:sameAs` as predicate, the process ends.

5.2 Dswget Design issues: an overview

In designing Dswget several issues, typical of the distributed systems, have been faced. Here we briefly report on the main of them without getting into too technical details.

In the Web of Data, a client in order to get information about a resource issues an HTTP GET request toward the HTTP server where the resource is hosted. In the standard case, the HTTP protocol offers a blocking semantics for its primitives, which means that once a request is issued the client has to wait for an answer or until a timeout. In Dswget, since engines exchange messages and data in a P2P fashion, a blocking semantics for communications would block the whole execution. To face this issue, specific asynchronous communication primitives and a *job delegation* mechanism have been implemented. With job delegation we mean that the sending Dswget engine delegates part of the execution and evaluation of a (sub)NAUTILOD expression to the receiving engine(s). In this respect, since a request, through the mechanism of *job delegation* is spread among multiple Dswget engines it is necessary to handle the ter-

mination of requests to avoid to keep consuming resources in an uncontrolled way. Dswget tackles this issue from two different perspectives:

(1) *Loop detection*: each Dswget engine keeps track, for each request, of each URI along with the state of the automaton on which it has been processed.

(2) *Termination*: this problem can be addressed by each Dswget engine which, for each request it receives informs the client that initially issued the request about the fact that it has operated on this request and whether it has delegated other Dswget engines. Then, the client can keep track of the list of the active engines on a particular request. The Dswget engine may additionally send back to the client the state of the automaton on which it is operating, thus enabling the client to know how far the execution is from a final state.

6. RELATED WORK

Many of the ideas underlying our proposal have been around in particular settings. We owe inspiration to several of them.

Navigation and specification languages of nodes in a graph have a deep research background. Nevertheless, most of its developments assume that data is stored in a central repository (e.g. Web query languages [9], *XPath*, navigational versions of SPARQL [22, 1]). They were inspiration for the navigational core of NAUTILOD.

Specification (and retrieval) of collections of sites was early addressed, and a good example is the well known tool *wget*. Besides being non-declarative, it is restricted to almost purely syntactic features. At semantic level, Hart et al. [18] proposed LDSpider, a crawler for the Web of Data able to retrieve RDF data by following RDF links according to different crawling strategies. They have little flexibility and are not declarative. The execution philosophy of *wget* was a source of inspiration for the incorporation of actions into NAUTILOD and to the design of *swget*.

Distributed data management has been explored and implemented by P2P and similar approaches [26]. For RDF, RDFPeers [5] and YARS2 uses P2P to answer RDF queries. Systems for distributed query processing on the Web have also been devised, e.g. DIASPORA [24]. Our distributed version of *swget* borrows some ideas from these approaches.

Finally, it is important to stress the fact that there is a

solid body of work on query processing and navigation on the Web of Data. Three lines of research can be identified:

(1) Load the *desired* data into a single RDF store (by *crawling* the LOD or some sub-portions) and process queries in a centralized way. There is a large list of Triple Stores [17]. There have been also developments in indexing techniques for semantic data. Swoogle [8], Sindice [20] and Watson [6] among the most successful. Recently, Hart et al. [13] proposed an approximate index structure for summarizing the content of Linked Data sources.

(2) Process the queries in a distributed form by using a federated query processor. DARQ [23] and FedX [25] provide mechanisms for transparently query answering on multiple query services. The query is split into sub-queries that are forwarded to the individual data sources and their result processed together. An evaluation of federated query approaches can be found in [12].

(3) Extend SPARQL with navigational features. The SERVICE feature of SPARQL 1.1 and proposals like the one of Hartig et al. [15] extend the scope of SPARQL queries with navigational features [15, 14]. The system SQUIN, based on *link-traversal*, a query execution paradigm that discover on the fly data sources relevant for the query, permits to automatically navigate to other sources while executing a query.

As it can be seen, our approach has a different departure point: it focuses on *navigational* functionalities, thus departing from querying as in (2); emphasizes specification of autonomous distributed sources, as opposed to (1); uses SPARQL querying to enhance navigation, while (3) proceeds in the reverse direction; and incorporates actions that in some sense generalize procedures implicit in the evaluation over the Web (*e.g.*, “get data” in crawlers and “return data” in query languages).

7. CONCLUSIONS AND FUTURE WORK

We presented a language to navigate, specify fragments and perform actions on the Web of Data. It explicitly exploits the semantics of the data “stored” at each URI. We implemented it in a centralized setting to run over real-world data, namely the LOD network, showing the benefits it can bring. We also developed a distributed version as proof-of-concept of its feasibility, potentialities and challenges.

The most important conclusion we can draw from this research and development is that the semantics given by RDF specifications *can be used* with profit to navigate, specify places and actions on the Web of Data. We presented a language that can be used as the basis for the development of agents that get data; navigate and report while navigating; and that can work immediately over LOD.

A second relevant finding we would like to report here, are the limitations found to take full advantage of the language and tools we developed. They refer essentially to (1) lack of standards in the sites regarding the dereferencing of data; (2) lack of standard RDF metadata regarding properties of the sites themselves (*e.g.*, provenance, summary of contents, etc.); (3) weak infrastructure to host delegation of execution and evaluation (of the language) to permit distribution. Tackling these issues can be considered as our wish list to leverage the Web of Data.

8. REFERENCES

- [1] F. Alkhateeb, J.-F. Baget, and J. Euzenat. Extending SPARQL with Regular Expression Patterns (for querying RDF). *J. Web Sem.*, 7(2):57–73, 2009.
- [2] T. Berners-Lee. Linked data design issues.
- [3] C. Bizer, T. Heath, and T. Berners-Lee. Linked Data - The Story So Far. *IJSWIS*, 5(3):1–22, 2009.
- [4] S. Brin and L. Page. The Anatomy of a Large-Scale Hypertextual Web Search Engine. *Computer Networks and ISDN Systems*, 30(1-7):107–117, 1998.
- [5] M. Cai and M. Frank. RDFPeers: a Scalable Distributed RDF Repository based on A Structured Peer-to-Peer Network. In *WWW*, 2004.
- [6] M. d’Aquin and E. Motta. Watson, more than a Semantic Web Search Engine. *Semantic Web*, 2(1):55–63, 2011.
- [7] S. J. DeRose, E. Maler, D. Orchard, and N. Walsh. XML Linking Language (XLink) Version 1.1. Technical report, 3 2010.
- [8] L. Ding, T. Finin, A. Joshi, R. Pan, R. S. Cost, Y. Peng, P. Reddivari, V. C. Doshi, and J. Sachs. Swoogle: A Search and Metadata Engine for the Semantic Web. In *CIKM*, 2004.
- [9] D. Florescu, A. Levy, and A. Mendelzon. Database techniques for the World-Wide Web: a survey. *SIGMOD Rec.*, 27:59–74, 1998.
- [10] G. Gottlob, C Koch, and R. Pichler. The Complexity of XPath Query Evaluation. In *PODS*, 2003.
- [11] C. Gutierrez, C. A. Hurtado, A. O. Mendelzon, and J. Pérez. Foundations of Semantic Web Databases. *J. Comput. Syst. Sci.*, 77(3):520–541, 2011.
- [12] P. Haase, T. Mathäb, and M. Ziller. An Evaluation of Approaches to Federated Query Processing over Linked Data. In *I-SEMANTICS*, 2010.
- [13] A. Harth, K. Hose, M. Karnstedt, A. Polleres, K. Sattler, and J. Umbrich. Data Summaries for On-demand Queries over Linked Data. In *WWW*.
- [14] O. Hartig. Zero-Knowledge Query Planning for an Iterator Implementation of Link Traversal Based Query Execution. In *ESWC*, 2011.
- [15] O. Hartig, C. Bizer, and J.-C. Freytag. Executing SPARQL Queries over the Web of Linked Data. In *ISWC*, 2009.
- [16] T. Heath and C. Bizer. *Linked Data: Evolving the Web into a Global Data Space*. Morgan & Claypool, 2011.
- [17] K. Hose, R. Schenkel, M. Theobald, and G. Weikum. Database Foundations for Scalable RDF Processing. In *Reasoning Web*, volume 6848, pages 202–249, 2011.
- [18] R. Isele, A. Harth, J. Umbrich, and C. Bizer. LDspider: An open-source crawling framework for the Web of Linked Data. In *Poster - ISWC*, 2010.
- [19] A. O. Mendelzon, G. A. Mihaila, and T. Milo. Querying the World Wide Web. *Int. J. on Digital Libraries*, 1(1):54–67, 1997.
- [20] E. Oren, R. Delbru, M. Catasta, R. Cyganiak, H. Stenzhorn, and G. Tummarello. Sindice.com: A document-oriented lookup index for open linked data. *Int. J. of Metad., Semant. and Ontolog.*, 3(1), 2008.
- [21] J. Pérez, M. Arenas, and C. Gutierrez. Semantics and Complexity of SPARQL. *ACM TODS*, 34(3), 2009.
- [22] J. Pérez, M. Arenas, and C. Gutierrez. nSPARQL: A Navigational Language for RDF. *JWS*, 8(4), 2010.
- [23] B. Quilitz and U. Leser. Querying Distributed RDF Data Sources with SPARQL. In *ESWC*, 2008.
- [24] M. Ramanath and J. R. Haritsa. DIASPORA: A Highly Distributed Web-Query Processing System. *World Wide Web*, 3(2):111–124, 2000.
- [25] A. Schwarte, P. Haase, K. Hose, R. Schenkel, and M. Schmidt. FedX: Optimization Techniques for Federated Query Processing on Linked Data. In *ISWC*, 2011.
- [26] P. Valduriez and E. Pacitti. Data Management in Large-Scale P2P Systems. In *VECPAR*, 2004.
- [27] P. Wadler. Two semantics for XPath, 1999. <http://www.cs.bell-labs.com/who/wadler/topics/xml.html>.