

ClioPatria: A SWI-Prolog Infrastructure for the Semantic Web

Editor(s): Axel Polleres, Vienna University of Economics and Business, Austria

Solicited review(s): Christoph Redl, Vienna TU, Austria; Alessandra Mileo, DERI, Ireland

Jan Wielemaker^a Wouter Beek^a Michiel Hildebrand^b Jacco van Ossenbruggen^b

^a *VU University Amsterdam, De Boelelaan 1081a, 1081 HV Amsterdam, The Netherlands,*
e-mail: {J.Wielemaker,W.G.J.Beek}@vu.nl

^b *CWI, Science Park 123, 1098 XG Amsterdam, The Netherlands,*
e-mail: {M.Hildebrand,Jacco.van.Ossenbruggen}@cwi.nl

Abstract. ClioPatria is a comprehensive semantic web development framework based on SWI-Prolog. SWI-Prolog provides an efficient C-based main-memory RDF store that is designed to cooperate naturally and efficiently with Prolog, realizing a flexible RDF-based environment for rule based programming. ClioPatria extends this core with a SPARQL and LOD server, an extensible web frontend to manage the server, browse the data, query the data using SPARQL and Prolog and a Git-based plugin manager. The ability to query RDF using Prolog provides query composition and smooth integration with application logic. ClioPatria is primarily positioned as a prototyping platform for exploring novel ways of reasoning with RDF data. It has been used in several research projects in order to perform tasks such as data integration and enrichment and semantic search.

Keywords: Triple Store, Logic Programming, Semantic Web Framework

1. Introduction

Many Semantic Web applications converge to a three-tier architecture¹—storage, application logic, presentation—that is popular in the relational database world, where SQL is replaced by SPARQL. SPARQL servers provide a varying and sometimes configurable level of entailment reasoning to support applications in data integration (e.g., `owl:sameAs`, `rdfs:subPropertyOf`) and hierarchical reasoning (e.g., `rdfs:subPropertyOf` and `rdfs:subClassOf`). The application logic is generally expressed in an (object-oriented) general purpose programming language.

Although we acknowledge there is a need for such architectures, because of their proven robustness, scalability and familiarity for the software industry, we

claim that these architectures are not suitable to explore the full potential of the ideas behind the Semantic Web for two reasons. Firstly, the separation between storage (SPARQL) and application logic (general purpose language) often causes a single computation to be expressed in a mixture of these two languages. This mixture harms readability as well as performance. Secondly, most languages used for the logic tier do not support higher level rule-based reasoning easily and suffer from the object-relational impedance mismatch [14], which makes expressing the application logic cumbersome.

We have developed ClioPatria to explore applications of the Semantic Web that are not well supported by the three-tier architecture. Its first application [25, 42] became the winner of the ISWC-2006 Semantic Web challenge. This application provides search over multiple museum collections while clustering the results based on their relation to the search term. The core of the application is an RDF path finding heuristic

¹<https://msdn.microsoft.com/en-us/library/ee658109.aspx>

and an algorithm for clustering paths based on the class and property hierarchies. Such techniques could not be expressed in SPARQL 1.0. Although SPARQL 1.1 primitives such as path-expressions allow for pushing a little more of the application logic to the query language, SPARQL-based solutions remain a mixture of SPARQL queries and client code that is hard to maintain and performs poorly due to the large number of SPARQL queries required to perform the search.

ClioPatria is a Prolog-based RDF framework that aims primarily at *prototyping*. It particularly facilitates prototyping due to the following features:

- A dedicated RDF store that is implemented in the C-language and designed to interface to Prolog naturally, providing a stable and efficient core storage and rule engine.
- Libraries that support all standard RDF interchange formats, the SPARQL 1.1 query language and LOD services, allowing to interface to other tools from the Linked Data community,
- The Prolog language provides a stable rule formalism and facilitates prototyping through live recompilation of the program under development, strong reflective and meta-programming support and the ability to ‘retry’ (go back to a previous point in the execution after undoing all assignments and execute the same code again) goals in the debugger that did not produce the expected result and thus inspect the execution at a more detailed level. [44]
- An integrated web development framework allows for serving an entire RDF-based web application from a single executable.

In this article, ClioPatria refers to both the ClioPatria application and the underlying SWI-Prolog libraries, notably the *semweb*, *RDF*, *pengines* and *http* packages. ClioPatria and these libraries have been developed together, where stable and general code typically became part of the SWI-Prolog libraries, while ClioPatria combines the libraries into a coherent and comprehensive application framework.

We would like to position ClioPatria primarily as an RDF *library* that is tightly connected to a programming language (Prolog). This positions it relatively close to Java-based systems such as Sesame, Jena, the Python-based rdflib and the Common Lisp-based AllegroGraph (see Section 3). It is positioned at a much greater distance from (often) disk-based SPARQL engines such as 4store and Virtuoso.

This document is organised as follows: before getting to related work in Section 3 we make the reader familiar with ClioPatria’s concepts, features and design in Section 2. Application areas are described in Section 4, where we show ClioPatria’s impact. In Section 5 we discuss the current limitations and our plans for the future. Section 6 concludes.

2. Introducing ClioPatria

2.1. Prolog and RDF

The central feature of ClioPatria is to use Prolog as rule and ‘glue’ language. RDF triples are naturally expressed in Prolog by the relation `rdf(Subject, Predicate, Object)`, where URIs are represented by Prolog atoms (interned strings) and literals are represented by a Prolog term that expresses the lexical form and the datatype or language tag. First, we give some examples of triples represented in Prolog taken from the RDF 1.1 primer document:²

```

rdf('http://example.org/bob#me',
    'http://...rdf-syntax-ns#type',
    'http://xmlns.com/foaf/0.1/Person').
rdf('http://example.org/bob#me',
    'http://xmlns.com/foaf/0.1/knows',
    'http://example.org/alice#me').
rdf('http://example.org/bob#me',
    'http://schema.org/birthDate',
    literal(type('http://...#date',
                '1990-07-04'))).

```

Like SPARQL and Turtle, the Prolog RDF layer allows for prefix notation. Prefixes are declared using `rdf_register_prefix(Prefix, URI)`, after which URIs can be written as *prefix:localname*. We will use this notation in the remainder of this article, using the conventional prefixes as defined on `http://prefix.cc/`.

Given an RDF database, Prolog allows us to write graph patterns as they appear in e.g., SPARQL as a conjunction of `rdf/3` relations. For example to get the birth dates of people somebody knows we can write:

²<http://www.w3.org/TR/2014/NOTE-rdf11-primer-20140624/#section-n-triples>

```
known_with_birthdate(Me, He, Date) :-
  rdf(Me, foaf:knows, He),
  rdf(He, schema:birthDate, Date).
```

The above code corresponds to the following SPARQL query:

```
SELECT ?me ?he ?date WHERE
{ ?me foaf:knows ?he .
  ?he schema:birthDate ?date
}
```

We claim, however, that the `known_with_birthdate/3` version in Prolog has many advantages over the SPARQL version, which we will explain below.

First, because the Prolog version has an explicit name, it is easy to reuse it in other queries. This allows developers to incrementally build more sophisticated functionality on top of existing, proven and tested building blocks. This is an effective way to avoid the large and complex SPARQL queries often found in other Semantic Web applications.

A second advantage is that the three arguments to `known_with_birthdate/3` can be used both to query and to filter. For readers that are not familiar with Prolog: identifiers that start with a capital letter are *variables*, while identifiers that start with a lower case letter or appear within single quotes are *constants*, called *atoms* in Prolog.

The interactive session below illustrates the use of `known_with_birthdate/3` to query the triple store. Note that the predicate generates all solutions one by one (the semicolon (;) is typed by the user to ask for the next solution).

```
?- known_with_birthdate(Me, He, Date).
Me = 'http://example.org/bob#me'
He = 'http://example.org/alice#me'
Date = literal(type(xsd:date,
                  '1990-07-04') ;
Me = 'http://example.org/bob#me'
He = 'http://example.org/john#me'
Date = literal(type(xsd:date,
                  '1992-09-06') ;
...

```

Instead of passing variables, we can also pass concrete values as arguments to the predicate. This is logically equivalent to filtering all results to those where the argument has the specified value. Calling

`known_with_birthdate('http://.../bob#me', He, BirthDate)` results in the first call to `rdf/3` using the combined subject-predicate index instead of just the predicate index, producing the reduced solution set efficiently. Calling `known_with_birthdate/3` with only a given *BirthDate* also produces the correct results, but inefficiently. This is discussed below.

Combined with operations on literals and Prolog's aggregation primitives, the above apparatus allows any SPARQL query to be expressed naturally. And, it can do more. Firstly, the above `known_with_birthdate/3` rule can be used as an abstraction of the two `rdf/3` statements and can be embedded in more complex graph patterns. This ability to *compose* queries facilitates exploring RDF data by assembling a library of simple graph patterns that can be named, documented and shared. Secondly, external sources such as relational databases or arbitrary Prolog relations can naturally be mixed with the RDF query. Thirdly, more powerful control structures can deal with recursion, constraint satisfaction, etc.

Prolog's *SLD resolution*, basically depth-first traversal of the search space, solves the stated queries under the provision that the user takes measures to avoid unbounded recursion.³ The performance of the SLD resolution strategy is very sensitive to proper ordering of the graph patterns. Prolog's *reflective* capabilities that allow Prolog to inspect and rewrite Prolog programs can solve this. The predicate `rdf_optimise(PatternIn, PatternOut)` rearranges a graph pattern based on statistics maintained by the core triple store [38].

2.1.1. Mapping more SPARQL concepts

Above, we described plain triple matching and simple conjunctions. Other SPARQL constructs are typically mapped to Prolog control primitives or higher-order (meta-)predicates. For example, UNION maps to Prolog disjunction (`QueryA; QueryB`) and OPTIONAL maps the the control structure (`Query *->true; true`). Aggregation, ordering and limiting the number of solutions are provided by the ISO Prolog aggregation predicates `bagof/3`, `setof/3` and `findall/3`, the de-facto standard library aggregate which provides *count*, *sum*, *average*, etc. and the SWI-Prolog library `solution_sequences` which provides *distinct*,

³Some Prolog systems, such as XSB and YAP, can detect and avoid unbounded recursion automatically using a technique called *tabling* [29].

limit, *order_by*, etc. The latter library was developed specifically to allow for a straightforward translation of many SPARQL queries to Prolog.

Below we give an example from the *Talk of Europe* project, which translated the European parliament debates to linked data that is accessible through ClioPatria.⁴ The query returns the first 10 tuples {Lang, Count} from the relation representing the number of speeches spoken in language Lang, ordered decreasingly by Count. The predicate `speech_language/2` defines a binary relation between a speech and the language in which it was spoken.

```
?- limit(10,
    order_by([desc(Count)],
    aggregate(count, Speech,
    speech_language(Speech, Lang),
    Count))).
```

The current library does not support SPARQL *property paths* explicitly. Support could be added by means of Prolog’s macro expansion facility implemented by `goal_expansion/2`. The code below provides a simplified partial implementation that illustrates the approach.

```
goal_expansion(rdf(S, P1/P2, O),
    ( rdf(S, P1, X),
    rdf(X, P2, O) )).
```

2.1.2. Expressiveness

Due to various restrictions, standardized query languages for SW data have limited expressivity. In practice, this requires the combination of server-side query execution and client-side post-processing of the results. There are several ways in which ClioPatria provides a more expressive programming environment.

Literal search in SPARQL is limited to the use of regular expressions and (numerical) conditions. ClioPatria couples its full-text literal index (Section 2.2) to Prolog’s strong background in Natural Language Processing (NLP) [41]. This allows more complicated operations such as stemming, phonetic matching, string distance calculation, and NLP parsing to be performed on literals. The alternative approach of using a general text indexing engine besides a triple store results in increased memory requirements and access times, and is difficult to synchronize under dynamic RDF triple assertions and retractions.

Even though SPARQL 1.1 is more powerful than version 1.0, there are still realistic queries that cannot be expressed in it [17]. For example, “which property paths connect a search string on literals to resources of type *Painting*”. For a fixed length, this can be expressed as follows in SPARQL (searching for *query* at distance 2):

```
SELECT ?l ?p1 ?r1 ?p2 ?r WHERE
{ ?r rdf:type art:Painting .
  ?r ?p1 ?r2 .
  ?r2 ?p2 ?l .
  FILTER regex(l, "query")
}
```

SPARQL does not provide the primitives to generalize this query for arbitrary path lengths, while this is easily expressed in Prolog:

```
literal_path_to_painting(Q, R, [L-P0|P]) :-
    rdf(R0, P0, literal(substring(Q), L)),
    path_to_painting(R0, R, P, []).

path_to_painting(R, R, [], _) :-
    rdf(R, rdf:type, art:'Painting').
path_to_painting(R0, R, [R0-P|Path], V) :-
    rdf(R0, P, R1),
    \+ memberchk(R1, V),
    path_to_painting(R1, R, Path, [R1|V]).
```

The above algorithm implements naive depth-first search. In the previously mentioned ISWC challenge application we implemented a best-first algorithm where the distance was based on predicate semantics (e.g., `owl:sameAs` does not change the distance) and predicate statistics (a predicate that links to more resources increments the distance more). It should be clear that SPARQL is not of much value for such problems.

2.1.3. Reasoning

Entailment reasoning can be implemented in several ways. ClioPatria provides a library that provides RDFS entailment primitives such as `rdfs_individual_of(?Resource, ?Class)`. This is the same as `{?resource rdf:type ?class}` after computing all triples entailed by the RDFS semantics. In addition, ClioPatria provides *entailment modules*. An entailment module is a Prolog module that extends the relation `rdf/3` which initially expresses only the plain triples. For example, the transitive semantics of `rdfs:subClassOf` can be expressed in an entailment module with the rule below, exploiting

⁴<http://www.talkofeurope.eu>

the transitive closure built-in `rdf_reachable/3`. The goal `rdf_reachable(S, rdfs:subClassOf, O)` implements the SPARQL expression `{?S rdf:type/rdfs:subClassOf* ?O}` using breath-first search.

```
rdf(S, rdfs:subClassOf, O):-
    rdf_reachable(S, rdfs:subClassOf, O).
```

Besides entailment models for standardized entailment regimes such as RDFS and OWL, it is just as easy to write an entailment module with custom, domain-specific rules. This provides similar functionality as *functional properties* in Jena, but where Jena restricts functional properties to compute a single object for a given subject and predicate,⁵ entailment rules can add zero or more *virtual triples* based on {S,P,O} pattern. According to the following example, persons younger than 18 are children:

```
rdf(Person, rdf:type, dbo:'Child'):-
    rdf(Person, dbo:age, Age),
    xsd_smaller(Age, 18).
```

Both forward and backward reasoning can be naturally expressed in Prolog. In many scenarios backward reasoning is to be preferred, because it reduces the memory footprint, avoids the problem of keeping inferred triples consistent after database modifications, and allows a single application to use different entailment definitions. Some deductions, however, are expensive to implement using backward reasoning. For example, finding all `rdf:type` triples that are entailed by RDFS domain and range restrictions requires a considerable amount of search. In such cases one can opt for forward chaining instead. Forward chaining must be coded explicitly. The fragment below adds derived type triples to the graph `rdfs`. The predicate `rdf_generation/1` can be used to verify whether the fixed point has been reached.

```
forall(( rdf(_,P,O),
         rdf(P, rdfs:domain, T)
       ),
       rdf_assert(O, rdf:type, T, rdfs)).
```

⁵<https://jena.apache.org/documentation/javadoc/jena/com/hp/hpl/jena/ontology/FunctionalProperty.html>

2.2. The triple store

Although triples could be expressed simply using a Prolog *dynamic predicate*, we have chosen for a low-level C implementation. This allows us to exploit the restricted form of this relation (no rule *body*, URIs as subject and predicate and only restricted ground terms as object) to reduce memory requirements and improve indexing. The underlying store is a memory-based quad-store that provides 9 hash-based indexes (s=subject, p=predicate, o=object, g=graph): s, p, o, sp, po, spo, g, sg and pg. These indexes are created and resized lazily. In addition, a closure of the `rdfs:subPropertyOf` relation is maintained to speed up queries that rely on subproperty entailment. The entailment of `rdfs:subPropertyOf` is treated special because we consider this property the key property for data alignment on the web of data and the space requirements are modest due to the typically relatively low number of properties in a dataset. These implementation choices have been motivated in [43]. An ordered index of all RDF literals is maintained to facilitate prefix or literal range queries. Optionally, the store can maintain a full-text index of all literals that provides fast search for literals that contain specified tokens, optionally extended with stemming and methaphone ('sounds-as').

The triple store has strong support for dynamic changes. Dynamic changes fit well with Prolog's bias towards backward chained reasoning (Section 2.1.3). Modifications to RDF data follow the Prolog *logical update view*, which implies that the results of queries are not affected by changes to the triple store that are made during query execution [40]. The triple store is designed to cooperate with the multi-threaded Prolog core. Read operations are implemented using lock-free techniques. Write operations use (short held) locks. Multiple updates may be combined into atomic *transactions* that can be nested. Prolog *failure* or a Prolog exception causes the transaction to be discarded. The same technology supports *snapshots*, which allows a Prolog goal to make isolated and temporary changes to the RDF store (Section 4.6.2).

The triple store optionally provides reliable persistency based on a *journal* file and fast save/load based on an optimized binary representation for RDF data.

Table 1 gives some core metrics of the RDF store. The store has been used with up to 200 million triples. Additional information can be found in [39].

Table 1

Core metrics of the SWI-Prolog RDF store. The values are based on the 64-bit version running Ubuntu 14.04 on an Intel i7-2600K@3.40Ghz CPU.

Metric	Value	Explanation
Memory	300 bytes/triple	Rough approximation of the storage needed per triple, including indexing. The value depends on the number of distinct URIs and literals as well as the length of URIs and literals.
Add triple (I)	3.5 μ s	Add a triple with a URI object that is computed in Prolog, where the URI is already interned as a Prolog atom.
Add triple (II)	7.5 μ s	Add a triple with a literal object that is computed in Prolog. This includes the time needed to intern the literal into the ordered literal store of the RDF database.
First triple	1.4 μ s	Get the first triple from an <code>rdf/3</code> query. This is almost independent from the number of triples in the store and whether the arguments are known.
Next Triple	0.6 μ s	Get the next triple from a <code>rdf/3</code> query.

2.3. Semantic web standards

The primary asset of the RDF initiative is a common data model and languages to operate on this model. ClioPatria contains modules for reading/writing the most standard RDF serialization formats (i.e., N-Quads, N-Triples, RDFa, RDF/XML, TriG and Turtle). The parsers and generators have been stress-tested extensively within the LOD Laundromat project (Section 4.3).

ClioPatria supports SPARQL 1.1. SPARQL queries are compiled into Prolog control structures. The resulting structure is optimized (Section 2.1) before execution. All SPARQL functions are implemented in a SPARQL runtime support library. The current implementation applies almost all these functions after performing the (optimised) graph pattern query, e.g. ignoring opportunities to map regular expression prefix searches to its ordered literal index. Interactive exploration is supported by the YASGUI query editor [24].

Support for Linked Data is available as a configurable option. Based on content negotiation, ClioPatria dereferences resources with machine-readable RDF or a human-readable HTML page. ClioPatria can serve URIs directly but also supports handling of URI redirection, e.g., through using `http://purl.org`. By default, the machine-readable dereference of a resource is the concise bounded graph of that resource. Alternatives can be selected or programmed as Prolog rules. It is possible to install custom HTML rendering of a resource, for example to define the page layout, suppress properties, include the rendering of related resources or render images.

2.4. Plugins & community extensions

A CPACK (or ClioPatria package) extends ClioPatria by providing additional Prolog libraries and/or web services. Packs are available from the ClioPatria site⁶ and can be installed from the Prolog shell. A CPACK consists of a directory hierarchy that provides Prolog libraries, configuration files and web resources such as CSS and JavaScript files. It is accompanied with RDF meta data that explains the type, purpose and authors of the package. In addition, the RDF description may specify dependencies. The CPACK is made available as a Git repository. If a CPACK is submitted, its Prolog code is analysed for dependencies on and conflicts with known CPACKs. Installing a CPACK installs all dependencies, adds the new directories to ClioPatria's file search paths and optionally loads them into the system based on the default configuration information provided. Due to the strong dependency on code analysis before accepting a package, packages typically coexist flawlessly without human intervention. The drawback is that a package must follow quite strict guidelines with respect to code quality and organization.⁷ For example, where (SWI-)Prolog loads and runs incomplete programs, the package analyzer demands that all referenced files and predicates exists, either in the plain Prolog system, plain ClioPatria or one of the other packages.

⁶<http://cliopatria.swi-prolog.org>

⁷<http://cliopatria.swi-prolog.org/help/>

2.5. Application development

Initially, ClioPatria applications are either developed as CPACKs (Section 2.4) or as independent Prolog source files that are loaded into ClioPatria. Recent versions provide a web-based alternative based on *Pengines*, or Prolog Engines [15].⁸ The Penguin infrastructure allows for embedding a Prolog program in a web page using a `<script type="text/x-prolog">` element and pose queries to this program from within JavaScript. If a query is submitted, both the query and program from the `<script>` elements are sent to the server using an HTTP POST request. The server loads the code into a sandboxed temporary Prolog module and executes the query. The client receives the results as JSON objects. Similar to the Prolog toplevel, the client can ask for additional solutions until it is satisfied.

This architecture allows web applications to be written without loading server extensions into ClioPatria, while the user can still profit from the advantages described in Section 2.1 such as query composition, data integration and the expressivity of the full Prolog language.

RDF data may be explored and suitable query fragments may be assembled using the Penguins-based web application called SWISH. SWISH is an acronym that stands for “SWI-Prolog for SHaring” or “SWI-Prolog SHell”. It provides a web-based code editor with code highlighting and template-based completion, a query editor with the same features and a result area. Programs can be saved, tagged and searched for to facilitate community-based management of code fragments.

3. Related work

Triples stores with reasoning capabilities bear some similarity to deductive databases such as XSB [28]. Even though Datalog systems are not directly applicable to SW data, [1] uses Datalog to extend the SPARQL 1.1 language with more powerful navigational patterns. Despite strong theoretical similarities, only few Logic Programming systems exist that specialize in storing and querying SW data. One such a system is GiaBATA [13] which allows queries to be performed against named graphs and under par-

tial RDFS and OWL entailment. It facilitates the implementation of custom reasoning capabilities through the specification of rules. Its implementation consists of a rewriting of SPARQL queries into Datalog programs. Instead of Prolog, it uses Answer Set Programming (ASP) as the embedded programming language. Where GiaBATA aims first of all at reasoning, concentrating on SPARQL, ClioPatria aims at providing a comprehensive semantic web platform where Prolog is the primary query language which can also be accessed through the Penguins API (see section 2.5). A similar approach is taken in the Smart-M3 platform [18], which is also based on ASP.

Euler [22] is an inference engine that uses hybrid reasoning and Euler path detection to avoid loops. As such, it is not intended to be used as a large-scale triple store, although it may use a triple store for input. Even though Euler runs on YAP and SWI-Prolog it is intended to be programmed in terms of Notation3 rules. Euler provides advanced reasoning paradigms such as abduction and induction besides deduction.

Several Semantic Web programming frameworks exist. Most frameworks are within the imperative and/or object-oriented paradigm, e.g., Jena [10] and Sesame [4] that are both based on Java and `rdflib`⁹ is based on Python. Querying in these libraries typically consists of creating an *iterator* from either a SPARQL query or a simple triple pattern. For example, `rdflib` defines `triples()`. This iterator can be used to find all people someone knows at distance two as follows:

```
for s,p,o in \
    g.triples((me,FOAF.knows,None)):
    for x,y,z in \
        g.triples((o,FOAF.knows,None)):
```

In Prolog, non-determinism avoids the need for explicit iterators and thus the fragment below suffices. Also note that this Prolog fragment is a non-deterministic goal and can thus be combined with other goals in a conjunction or disjunction. The Python loop above is *not* an iterator that can be reused ‘as is’.

```
rdf(X, foaf:knows, Y),
rdf(Y, foaf:knows, Z),
```

AllegroGraph¹⁰ is a commercial high-performance, persistent graph database. It combines in-memory with

⁸Pengines are online at <http://pengines.swi-prolog.org>

⁹<https://github.com/RDFLib/rdflib>

¹⁰<http://franz.com/agraph/allegrograph/>

disk-based storage, which allows it to be very scalable (billions of quadruples). AllegroGraph supports SPARQL, RDFS++ and Prolog reasoning. The DB-engines page on AllegroGraph¹¹ describes AllegroGraph first of all as an *integration framework*. The stress on integration is similar to ClioPatria. An obvious difference is that AllegroGraph’s principal ‘glue’ language is Common Lisp and the Prolog system (thus) uses Lisp syntax rather than Prolog syntax. We consider their inclusion of “Prolog reasoning for client applications” a strong support for ClioPatria’s architecture. Both based on a high-level declarative programming language, neither of the systems is specifically focussed on a particular application domain. A side-by-side comparison of features is hard to carry out and considered outside the scope of this document.

4. Application areas

In this section we briefly describe a number of applications based on ClioPatria or SWI-Prolog’s triple store. The applications are collected from our academic network, the mailinglist, literature research and searches on GitHub. We organised the applications in *application areas*, based on the way they (primarily) use ClioPatria’s features. The section may be read both as a proof of impact and as a list of examples of how ClioPatria can be deployed.

All mentioned projects are academic projects and, as we often depend on publications to become aware of a project, most projects are finished. Some applications survived after project completion as a CPACK, e.g., XMLRDF and Amalgame. Some are maintained as online services, e.g., Europeana Thought Lab. The KnowRob knowledge base is at the heart of an active research group on robotics. The LOD Laundromat project is an example of new research where ClioPatria plays a significant role.

4.1. Semantic search & faceted navigation

The semantic search application mentioned in Section 1 started the development of ClioPatria. This search engine is reused and extended in several projects, e.g., CHIP and Europeana Thought Lab. The core of the semantic search application is a graph search algorithm. ClioPatria’s tight integration of the

RDF library with the program language allowed for an intuitive way to implement the best-first search algorithm.

/facet is another early project using ClioPatria [12]. The application provides faceted navigation on arbitrary RDF datasets. It was used to explore collections of art objects [12], music (DBtune, [23]) and news images [31]. For an RDF dataset, */facet* precomputes per `rdfs:Class` the facets that are available for its instances. It also detects abstract facets that can be derived using the `rdfs:subPropertyOf` relations in the data. For example, in the news domain the detailed properties of the NewsML ontology were defined as subProperties of the *who*, *what*, *where* and *when* abstract properties [11]. In the user interface these abstract facets provide a convenient and understandable view over the data collection at the start of a browsing session. If needed, the facets corresponding to the specific properties can be activated to construct more fine-grained queries. In a similar fashion the application detects if the values within a single facet are hierarchically organized, e.g., by the `rdfs:subClassOf` or `skos:broader` relations, and precomputes the spanning tree of the values in this facets. In the interface, the user can then use the top level concepts of these hierarchies to navigate the collection. The facet detection and abstraction technique as well as computing the results for faceted queries are implemented as Prolog rules in ClioPatria. With the introduction of property paths in SPARQL 1.1 it has become feasible to express most of the computations neatly in SPARQL.

In [32], ClioPatria is used to develop a mobile cultural heritage guide that combines information from the LOD cloud and cultural heritage institutes in order to display information about points of interest on a mobile phone. In this project ClioPatria serves multiple purposes. Firstly, as a triple store with common SW functionality such as graph search, semantic crawling and reasoning. Secondly, domain-specific reasoning capabilities, in this case spatial reasoning, are integrated on top of the generic SW reasoning capabilities. Thirdly, since mobile devices may have limited connectivity and processing power, ClioPatria performs the computationally heavy task of facet classification on the server side. In this way only the relevant information needs to be transmitted to the client application.

In [6,5] ClioPatria is used to allow scholars to search for cultural heritage resources in the Digital Archives

¹¹<http://db-engines.com/en/system/AllegroGraph>

of Italian psychology¹² in terms of the CIDOC-CRM¹³ ontology.

4.2. Web application generator & platform

ClioPatria provides an integrated comprehensive web application platform, capable of dealing with all three tiers: storage, application logic and presentation. Most of the search applications from Section 4.1 are completely coded in ClioPatria.

As another example, Amalgame [36] provides a web-based interface for interactive vocabulary alignment. It is distributed as a ClioPatria CPACK (Section 2.4) and uses the triple store backend to store the RDF representations of the vocabularies that are aligned. It also stores an RDF representation of the alignment strategy that the user incrementally builds while interacting with the tool as well as the RDF representations of the alignment results. By using the triple store's persistency features, Amalgame exploits the advantages of a fast in-memory store with the possibility to continue tasks after server reboots or timed-out sessions. Amalgame has a flexible plugin-infrastructure that allows new or modified alignment algorithms to be inserted without restarting the web server or triple store. This is important in a research prototyping environment, where such changes happen frequently, and reloading large vocabularies after each change is too time consuming. Finally, Amalgame relies on the NLP extensions of the triple store to perform fast string matching between RDF literals (Section 2.1.2), for example using the built-in Porter stemmer and the *isub* [27] string distance matcher. The prefix index on literals is used to provide efficient auto-completion services. Such support is crucial as it allows the user to lookup terms interactively in order to correct alignment mistakes that were made by the system.

4.3. Web crawling & robustness

LOD Laundromat¹⁴ [3] is a system that performs a large, LOD cloud-scale data cleaning operation. It searches for files and endpoints that might contain RDF triples starting from seed points. First, possible compression and archiving (e.g., ZIP) is analysed

based on the SWI-Prolog binding to libarchive.¹⁵ Next, for each individual file, heuristics are used to detect whether the data contains RDF and —if so— using what serialization format. Finally, the data is parsed by the applicable RDF parser from ClioPatria while keeping statistics about errors and warnings. LOD Laundromat has processed hundreds of thousands of data documents, parsing tens of billions of triples. All data is disseminated into a canonical and standards-compliant format. This project shows that ClioPatria has sufficient robustness to support large-scale processing.

4.4. Data integration & enrichment

ClioPatria, due to its underlying triple store that supports fast changes to the RDF data and backward chaining based reasoning that does not suffer from maintaining derived entailed information compatible with changes to the triple set, has proven to be a suitable platform for combining RDF data with external data, as well as for generating and restructuring RDF data. In such data integration scenarios, Prolog is used as the rule language and 'glue' towards other data sources such as relational databases.

SWI-Prolog's RDF infrastructure has been used to convert WordNet to RDF [33]. The XMLRDF toolkit,¹⁶ written on top of ClioPatria, converts XML datadumps into RDF using a generic mapping, after which the RDF can be transformed into the desired shape by using rules that are translated into Prolog. XMLRDF has been used in the Europeana Digital Library project as well as for the conversion and enrichment of many Dutch cultural heritage datasets and the-sauri [9,8].

The Bliokit [20] system is a comprehensive system for querying and transforming data in bioinformatics. Bliokit's RDF support in terms of performance and expressiveness was evaluated by Lampa [16]. Lampa stresses the enhanced maintainability of Prolog queries due to reuse of query fragments represented as predicates (Figure 18 in [16]).

4.5. Wiki systems

PIWiki [21] is a semantic wiki implemented on top of ClioPatria. The authors claim that a crucial reason for using a Prolog based solution is the transparent in-

¹²Archivi Storici della Psicologia Italiana (ASPI), <http://www.aspi.unimib.it>

¹³<http://www.cidoc-crm.org/>

¹⁴<http://lodlaundromat.org/>

¹⁵<http://www.libarchive.org/>

¹⁶<http://semanticweb.cs.vu.nl/xmlrdf/>

tegration of various rule formalisms such as SWRL and XTT.

4.6. Reasoning

We collected some projects related to reasoning on top of the Prolog RDF store. Many projects about reasoning do not require the higher level ClioPatria features.

4.6.1. Robotics

The openEASE¹⁷ uses the SWI-Prolog RDF infrastructure for the KnowRob [30] knowledge base, which represents ontological knowledge in OWL using the SWI-Prolog RDF store. Dedicated reasoning components that perform for example spatial reasoning or visibility computations are linked to the OWL representation is *procedural attachments*. These attachments can both extend the OWL model with inferred information as well as pose additional constraints on the OWL descriptions. This example illustrates the value of a tight integration between the RDF description and a programming language that naturally extends the RDF relational model.

4.6.2. Counterfactual reasoning

*Actual-Causation*¹⁸ is a research project in which causal models can be created and simulated. According to a popular definition in analytic philosophy, causation is defined in terms of counterfactuals, i.e. conditions that would apply if the course of events had been different. Implementation-wise, causal models are collections of RDF statements that reside in a ClioPatria triple store. The snapshot feature of ClioPatria is used to make isolated and short-living copies of the database, one for each counterfactual. Within each snapshot specific facts are added/retracted/alterd in order to express the respective counterfactuals. Backward chaining is used to validate whether certain causal effects still follow under those different circumstances. In a triple store that does not support snapshots, this would have been implemented less elegantly, e.g., by inefficiently copying data between an extensive number of named graphs.

4.6.3. Spatial reasoning

In the Poseidon project [34] ClioPatria is used to reason about ship trajectories. The library that allows

for combining Linked Data with spatial reasoning has been made available as a SWI-Prolog package [35] and is an example of integrating domain-specific knowledge with SW reasoning facilities into a single clean paradigm.

4.6.4. Qualitative reasoning

WebQR [2] is another SW-based domain-specific knowledge system that allows physical systems to be qualitatively simulated. In WebQR a user creates a model whose components are SW resources. A ClioPatria-based caching infrastructure is used to retrieve knowledge statements about those components from the web. Traditional RDF(S) entailment, in combination with domain-specific reasoning rules, is used in order to automate modeling steps that would otherwise be performed by human modelers. For example, WebQR can decide whether a modeled resource can vary over time (=quantity) or not (=entity). Whether a derivation is possible depends on whether specific information has yet been cached, using ClioPatria's dynamic triple storage capabilities.

5. Future work

ClioPatria is being actively developed and its features are continuously expanded to satisfy project requirements. Below we give a list of recognised limitations.

1. The current system has limited scalability of its RDF store because it is purely memory based. The strong dynamic RDF support of the platform allows the local store to be used as a cache which relies on large external stores for background knowledge. We plan to provide an interface to *linked data fragments* [37]. Using the LOD Laundromat linked data fragments endpoint and appropriate caching techniques will provide web-scale querying.
2. Reasoning that requires recursion often puts the onus on the user to explicitly avoid cycles. This can be avoided by using *tabling*, a technique that is available in the Prolog systems such as XSB and YAP, but not yet in SWI-Prolog.
3. Standardized SW reasoning in ClioPatria is currently limited to RDFS++. For example, full-scale OWL-DL reasoning is not yet supported, although this should be relatively easy given the underlying architecture.

¹⁷<http://www.open-ease.org/getting-started/>

¹⁸<http://wouterbeek.github.io/>

4. ClioPatria implements SPARQL 1.1 query and update requests, but not yet SPARQL federation.
5. The RDF store currently represents literals as a tuple of type or language and lexical representation represented as a Prolog atom. For example, the literal representing the latitude of a resource is a term `type(Type,Atom)`, where `Type` is a float or decimal XML schema type and `Atom` contains the lexical representation of the number. A Prolog program dealing with this property requires a Prolog floating point number. This implies that the Prolog programmer needs to include an explicit conversion step. This step makes the code less readable and unnecessarily slow. We plan to provide smooth querying through the *value space*. The new API for this will include the Turtle based `String@lang` and `Value^^Type` syntax.
6. While Lampa [16] claims that queries written in Prolog often outperform competing SPARQL endpoints, poor integration of SPARQL literal operations into the execution plan does not yet expose this performance to the SPARQL endpoint of ClioPatria.

The current development concentrates on the web-based shared Prolog development environment SWISH (see section 2.5). SWISH offers collaborative development and maintenance of queries that resemble a wiki. This context gives priority to issues one and five mentioned above. We also plan to offer a SPARQL compliant result format for the Prolog query API. Synchronizing the result format simplifies reuse of existing client software as well as migrating applications between SPARQL and Prolog endpoints.

6. Conclusion

In this paper we have presented the ClioPatria Semantic Web toolkit. ClioPatria is based on SWI-Prolog and is tightly connected to an efficient main-memory RDF quad store. We have shown that the use of ClioPatria and its Logic Programming (LP) paradigm has several benefits as a prototyping environment for SW programming. Accessing RDF data using LP does not suffer from the object-relational impedance mismatch and LP allows expressions to be built from small reusable components. By providing a general purpose language closely connected to the triple store we can formulate more expressive queries in graphs as well as on literals. An emphasis on backward chaining makes LP suitable

for reasoning with dynamic RDF. This feature is reflected in the database by supporting transactions and isolated modifications in *snapshots*.

The toolkit is completed by standard SW functionality such as a SPARQL 1.1 endpoint, parsers and serializers for most RDF formats, a plugin infrastructure and a web frontend. The most important application areas for ClioPatria are currently data integration and enrichment and semantic search.

The recent introduction of Pengines (Prolog Engines) and SWISH (SWI-Prolog SHaring/SHell) allows the use of flexible Prolog rule-based reasoning without changing the software on the server. Pengines allow for remote execution of Prolog programs and can be accessed both from JavaScript to realise web applications and from Prolog to realise distributed computing.

Acknowledgements

The development of ClioPatria has been initiated under the Dutch MultimediaN project. Further development has taken place in the context of many different projects. This publication was supported by the Dutch national program COMMIT/

References

- [1] Marcelo Arenas, Georg Gottlob, and Andreas Pieris. Expressive languages for querying the semantic web. In Georg Gottlob and Jorge Pérez, editors, *Proceedings of the 8th Alberto Mendelzon Workshop on Foundations of Data Management, Cartagena de Indias, Colombia, June 4-6, 2014.*, volume 1189 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2014.
- [2] Wouter Beek, Sander Latour, and Stefan Schlobach. Webqr: Building a knowledge representation application on the semantic web. In Ruben Verborgh and Erik Mannens, editors, *Proceedings of the ISWC Developers Workshop 2014, co-located with the 13th International Semantic Web Conference (ISWC 2014), Riva del Garda, Italy, October 19, 2014.*, volume 1268 of *CEUR Workshop Proceedings*, pages 102–107. CEUR-WS.org, 2014.
- [3] Wouter Beek, Laurens Rietveld, Hamid R. Bazoobandi, Jan Wielemaker, and Stefan Schlobach. LOD laundromat: A uniform way of publishing other people’s dirty data. In Mika et al. [19], pages 213–228.
- [4] Jeen Broekstra, Arjohn Kampman, and Frank van Harmelen. Sesame: A generic architecture for storing and querying RDF and RDF schema. In Ian Horrocks and James A. Hendler, editors, *The Semantic Web - ISWC 2002, First International Semantic Web Conference, Sardinia, Italy, June 9-12, 2002, Proceedings*, volume 2342 of *Lecture Notes in Computer Science*, pages 54–68. Springer, 2002.
- [5] Claudio Cortese. Semantic search and browsing nell’ambito dei beni culturali. Master’s thesis, Bollettino del CILEA, 2011.

- [6] Claudio Cortese and Glauco Mantegari. Extending the digital archives of italian psychology with semantic data. In Livia Predoiu, Steffen Henniecke, Andreas Nürnberger, Annett Mitschick, and Seamus Ross, editors, *Proceedings of the 1st International Workshop on Semantic Digital Archives, Berlin, Germany, September 29, 2011*, volume 801 of *CEUR Workshop Proceedings*, pages 60–71. CEUR-WS.org, 2011.
- [7] Isabel F. Cruz, Stefan Decker, Dean Allemang, Chris Preist, Daniel Schwabe, Peter Mika, Michael Uschold, and Lora Aroyo, editors. *The Semantic Web - ISWC 2006, 5th International Semantic Web Conference, ISWC 2006, Athens, GA, USA, November 5-9, 2006, Proceedings*, volume 4273 of *Lecture Notes in Computer Science*. Springer, 2006.
- [8] Victor de Boer, Matthias van Rossum, Jurjen Leinenga, and Rik Hoekstra. Dutch ships and sailors linked data. In Mika et al. [19], pages 229–244.
- [9] Viktor de Boer, Jan Wielemaker, Judith van Gent, Michiel Hildebrand, Antoine Isaac, Jacco van Ossenbruggen, and Guus Schreiber. Supporting linked data production for cultural heritage institutes: The amsterdam museum case study. In Elena Simperl, Philipp Cimiano, Axel Polleres, Óscar Corcho, and Valentina Presutti, editors, *The Semantic Web: Research and Applications - 9th Extended Semantic Web Conference, ESWC 2012, Heraklion, Crete, Greece, May 27-31, 2012. Proceedings*, volume 7295 of *Lecture Notes in Computer Science*, pages 733–747. Springer, 2012.
- [10] Michael Grobe. Rdf, jena, sparql and the 'semantic web'. In Gail Farally-Semerad, Karen J. McRitchie, and Elizabeth Rugg, editors, *Proceedings of the ACM SIGUCCS Fall Conference on User Services 2009, St. Louis, Missouri, USA, October 11-14, 2009*, pages 131–138. ACM, 2009.
- [11] M. Hildebrand and J. R. van Ossenbruggen. Configuring semantic web interfaces by data mapping. In *Proceedings of the VISSW 2009 Workshop: Visual Interfaces to the Social and the Semantic Web*, Proceedings of the VISSW Workshop: Visual Interfaces to the Social and the Semantic Web, February 2009.
- [12] Michiel Hildebrand, Jacco van Ossenbruggen, and Lynda Hardman. /facet: A browser for heterogeneous semantic web repositories. In Cruz et al. [7], pages 272–285.
- [13] Giovambattista Ianni, Thomas Krennwallner, Alessandra Martello, and Axel Polleres. A rule system for querying persistent RDFS data. In Lora Aroyo, Paolo Traverso, Fabio Ciravegna, Philipp Cimiano, Tom Heath, Eero Hyvönen, Riichiro Mizoguchi, Eyal Oren, Marta Sabou, and Elena Paslaru Bontas Simperl, editors, *The Semantic Web: Research and Applications, 6th European Semantic Web Conference, ESWC 2009, Heraklion, Crete, Greece, May 31-June 4, 2009, Proceedings*, volume 5554 of *Lecture Notes in Computer Science*, pages 857–862. Springer, 2009.
- [14] Christopher Ireland, David Bowers, Michael Newton, and Kevin Waugh. A classification of object-relational impedance mismatch. In *Proceedings of the 2009 First International Conference on Advances in Databases, Knowledge, and Data Applications, DBKDA '09*, pages 36–43, Washington, DC, USA, 2009. IEEE Computer Society.
- [15] Torbjörn Lager and Jan Wielemaker. Penguins: Web logic programming made easy. *TPLP*, 14(4-5):539–552, 2014.
- [16] Samuel Lampa. SWI-Prolog as a Semantic Web tool for semantic querying in Bioclipse: Integration and performance benchmarking. Master's thesis, Uppsala University, 2010.
- [17] Leonid Libkin, Juan L. Reutter, and Domagoj Vrgoc. Trial for RDF: adapting graph query languages for RDF data. In Richard Hull and Wenfei Fan, editors, *Proceedings of the 32nd ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS 2013, New York, NY, USA - June 22 - 27, 2013*, pages 201–212. ACM, 2013.
- [18] Vesa Luukkala and Ilkka Niemelä. Enhancing a smart space with answer set programming. In Mike Dean, John Hall, Antonino Rotolo, and Said Tabet, editors, *Semantic Web Rules - International Symposium, RuleML 2010, Washington, DC, USA, October 21-23, 2010. Proceedings*, volume 6403 of *Lecture Notes in Computer Science*, pages 89–103. Springer, 2010.
- [19] Peter Mika, Tania Tudorache, Abraham Bernstein, Chris Welty, Craig A. Knoblock, Denny Vrandecic, Paul T. Groth, Natasha F. Noy, Krzysztof Janowicz, and Carole A. Goble, editors. *The Semantic Web - ISWC 2014 - 13th International Semantic Web Conference, Riva del Garda, Italy, October 19-23, 2014. Proceedings, Part I*, volume 8796 of *Lecture Notes in Computer Science*. Springer, 2014.
- [20] Chris Mungall. Experiences using logic programming in bioinformatics. In Patricia M. Hill and David Scott Warren, editors, *Logic Programming, 25th International Conference, ICLP 2009, Pasadena, CA, USA, July 14-17, 2009. Proceedings*, volume 5649 of *Lecture Notes in Computer Science*, pages 1–21. Springer, 2009.
- [21] Grzegorz J. Nalepa. Plwiki - A generic semantic wiki architecture. In Ngoc Thanh Nguyen, Ryszard Kowalczyk, and Shyi-Ming Chen, editors, *Computational Collective Intelligence. Semantic Web, Social Networks and Multiagent Systems, First International Conference, ICCCI 2009, Wroclaw, Poland, October 5-7, 2009. Proceedings*, volume 5796 of *Lecture Notes in Computer Science*, pages 345–356. Springer, 2009.
- [22] Pieterjan De Potter, Hans Cools, Kristof Depraetere, Giovanni Mels, Pedro Debevere, Jos De Roo, Csaba Huszka, Dirk Colaert, Erik Mannens, and Rik Van de Walle. Semantic patient information aggregation and medicinal decision support. *Computer Methods and Programs in Biomedicine*, 108(2):724–735, 2012.
- [23] Yves Raimond and Mark B. Sandler. A web of musical information. In Juan Pablo Bello, Elaine Chew, and Douglas Turnbull, editors, *ISMIR 2008, 9th International Conference on Music Information Retrieval, Drexel University, Philadelphia, PA, USA, September 14-18, 2008*, pages 263–268, 2008.
- [24] Laurens Rietveld and Rinke Hoekstra. YASGUI: not just another SPARQL client. In Philipp Cimiano, Miriam Fernández, Vanessa Lopez, Stefan Schlobach, and Johanna Völker, editors, *The Semantic Web: ESWC 2013 Satellite Events - ESWC 2013 Satellite Events, Montpellier, France, May 26-30, 2013, Revised Selected Papers*, volume 7955 of *Lecture Notes in Computer Science*, pages 78–86. Springer, 2013.
- [25] Guus Schreiber, Alia K. Amin, Mark van Assem, Viktor de Boer, Lynda Hardman, Michiel Hildebrand, Laura HOLLINK, Zhisheng Huang, Janneke van Kersen, Marco de Niet, Borys Omelayenko, Jacco van Ossenbruggen, Ronny Siebes, Jos Taekema, Jan Wielemaker, and Bob J. Wielinga. Multimedial e-culture demonstrator. In Cruz et al. [7], pages 951–958.
- [26] Amit P. Sheth, Steffen Staab, Mike Dean, Massimo Paolucci, Diana Maynard, Timothy W. Finin, and Krishnaprasad Thirunarayan, editors. *The Semantic Web - ISWC 2008, 7th International Semantic Web Conference, ISWC 2008, Karlsruhe, Germany, October 26-30, 2008. Proceedings*, volume 5318 of

- Lecture Notes in Computer Science*. Springer, 2008.
- [27] Giorgos Stoilos, Giorgos B. Stamou, and Stefanos D. Kollias. A string metric for ontology alignment. In Yolanda Gil, Enrico Motta, V. Richard Benjamins, and Mark A. Musen, editors, *The Semantic Web - ISWC 2005, 4th International Semantic Web Conference, ISWC 2005, Galway, Ireland, November 6-10, 2005, Proceedings*, volume 3729 of *Lecture Notes in Computer Science*, pages 624–637. Springer, 2005.
- [28] Terrance Swift and David S. Warren. XSB: Extending Prolog with tabled logic programming. *Theory and Practice of Logic Programming*, 12:157–187, 1 2012.
- [29] Terrance Swift and David Scott Warren. XSB: extending prolog with tabled logic programming. *TPLP*, 12(1-2):157–187, 2012.
- [30] Moritz Tenorth and Michael Beetz. KnowRob – A Knowledge Processing Infrastructure for Cognition-enabled Robots. *International Journal of Robotics Research (IJRR)*, 32(5):566 – 590, April 2013.
- [31] Raphaël Troncy. Bringing the IPTC news architecture into the semantic web. In Sheth et al. [26], pages 483–498.
- [32] Chris J. van Aart, Bob J. Wielinga, and Willem Robert van Hage. Mobile cultural heritage guide: Location-aware semantic search. In Philipp Cimiano and Helena Sofia Pinto, editors, *Knowledge Engineering and Management by the Masses - 17th International Conference, EKAW 2010, Lisbon, Portugal, October 11-15, 2010. Proceedings*, volume 6317 of *Lecture Notes in Computer Science*, pages 257–271. Springer, 2010.
- [33] Mark Van Assem, Aldo Gangemi, and Guus Schreiber. Conversion of WordNet to a standard RDF/OWL representation. In *Proceedings of the Fifth International Conference on Language Resources and Evaluation*, pages 237–242, 2006.
- [34] Willem Robert van Hage, Véronique Malaisé, Gerben de Vries, Guus Schreiber, and Maarten van Someren. Abstracting and reasoning over ship trajectories and web data with the Simple Event Model (SEM). *Multimedia Tools Appl.*, 57(1):175–197, 2012.
- [35] Willem Robert van Hage, Jan Wielemaker, and Guus Schreiber. The space package: Tight integration between space and semantics. *T. GIS*, 14(2):131–146, 2010.
- [36] Jacco van Ossenbruggen, Michiel Hildebrand, and Viktor de Boer. Interactive vocabulary alignment. In Stefan Gradmann, Francesca Borri, Carlo Meghini, and Heiko Schuldt, editors, *Research and Advanced Technology for Digital Libraries - International Conference on Theory and Practice of Digital Libraries, TPDL 2011, Berlin, Germany, September 26-28, 2011. Proceedings*, volume 6966 of *Lecture Notes in Computer Science*, pages 296–307. Springer, 2011.
- [37] Ruben Verborgh, Miel Vander Sande, Pieter Colpaert, Sam Coppens, Erik Mannens, and Rik Van de Walle. Web-scale querying through linked data fragments. In Christian Bizer, Tom Heath, Sören Auer, and Tim Berners-Lee, editors, *Proceedings of the Workshop on Linked Data on the Web co-located with the 23rd International World Wide Web Conference (WWW 2014), Seoul, Korea, April 8, 2014.*, volume 1184 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2014.
- [38] Jan Wielemaker. An optimised semantic web query language implementation in prolog. In Maurizio Gabbriellini and Gopal Gupta, editors, *Logic Programming, 21st International Conference, ICLP 2005, Sùges, Spain, October 2-5, 2005, Proceedings*, volume 3668 of *Lecture Notes in Computer Science*, pages 128–142. Springer, 2005.
- [39] Jan Wielemaker. *Logic programming for knowledge-intensive interactive applications*. PhD thesis, University of Amsterdam, 2009. <http://dare.uva.nl/en/record/300739>.
- [40] Jan Wielemaker. Extending the logical update view with transaction support. *CoRR*, abs/1301.7669, 2013.
- [41] Jan Wielemaker, Michiel Hildebrand, and Jacco van Ossenbruggen. Prolog as the fundament for applications on the semantic web. In Axel Polleres, David Pearce, Stijn Heymans, and Edna Ruckhaus, editors, *Proceedings of the ICLP'07 Workshop on Applications of Logic Programming to the Web, Semantic Web and Semantic Web Services, ALPSWS 2007, Porto, Portugal, September 13th, 2007*, volume 287 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2007.
- [42] Jan Wielemaker, Michiel Hildebrand, Jacco van Ossenbruggen, and Guus Schreiber. Thesaurus-based search in large heterogeneous collections. In Sheth et al. [26], pages 695–708.
- [43] Jan Wielemaker, Guus Schreiber, and Bob Wielinga. Prolog-based infrastructure for RDF: performance and scalability. In D. Fensel, K. Sycara, and J. Mylopoulos, editors, *The Semantic Web - Proceedings ISWC'03, Sanibel Island, Florida*, pages 644–658, Berlin, Germany, october 2003. Springer Verlag. LNCS 2870.
- [44] Jan Wielemaker, Tom Schrijvers, Markus Triska, and Torbjörn Lager. Swi-prolog. *TPLP*, 12(1-2):67–96, 2012.