# Semantic Web Technologies:
# From Theory to Standards

Axel Polleres

Digital Enterprise Research Institute, National University of Ireland, Galway
`firstname.lastname@deri.org`

**Abstract.** This paper summarises the evolution of W3C standards in the area of Semantic Web technologies, as well as gaps within these standards still to be filled in terms of standardisation. Moreover, we give a subjective survey of the most influential scientific works which have contributed to the development of these standards and to closing the gaps between them. The Semantic Web proves to become an interesting application field for Artificial Intelligence; we aim here at both giving an overview of own work in the area as well as providing an entry point for researchers interested in the foundations of Semantic Web standards and technologies.

## 1 Introduction – The Semantic Web Architecture

The Semantic Web is about to grow up. Over the last few years technologies and standards to build up the architecture of this next generation of the Web have matured and are being deployed on large scale in many live Web sites. The underlying technology stack of the Semantic Web consists of several standards endorsed by the World Wide Web consortium (W3C) that provide the formal underpinings of a machine-readable "Web of Data" [94]:

- A Uniform Exchange Syntax: the eXtensible Markup Language (XML)
- A Uniform Data Exchange Format: the Resource Description Framework (RDF)
- Ontologies: RDF Schema and the Web Ontology Language (OWL)
- Rules: the Rule interchange format (RIF)
- Query and Transformation Languages: XQuery, SPARQL

**The eXtensible Markup Language (XML)** Starting from the pure HTML Web which mainly facilitated the exchange of layout information for Web pages only, the introduction of the eXtensible Markup Language (XML) in its first edition in 1998 [19] meant a breakthrough for Web technologies. With XML as a uniform exchange syntax, any semi-structured data can be modeled as a tree. Along with available APIs, parsers and other tools, XML allows one to define various other Web languages besides HTML. XML nowadays is not only the basis for Web data, but also for Web services [45] and is used in many custom applications as a convenient data exchange syntax. Schema description languages such as XML Schema [112] can be used to define XML languages; expressive query and transformation languages such as XQuery [27] and XSLT [68] allow for querying specific parts of an XML tree, or for transforming one XML language into another.

**The Resource Description Framework (RDF)** The *Resource Description Framework* (RDF) – now around for over a decade already as well – is the basic data model for the Semantic Web. It is built upon one of the simplest structures for representing data: a directed labeled graph. An RDF graph is described by a set of triples of the form $\langle Subject\ Predicate\ Object \rangle$, also called *statements*, which represent the edges of this graph. Anonymous nodes in this graph – so called-blank nodes, akin to existential variables – allow one to also model incomplete information. RDF's flat graph-like representation has the advantage of abstracting away from the data schema, and thus promises to allow for easier integration than customised XML data in different XML dialects: whereas the integration of different XML languages requires the transformation between different tree structures using transformation languages such as XSLT [68] or XQuery [27], different RDF graphs can simply be stored and queried alongside one another, and as soon as they share common nodes, form a joint graph upon a simple merge operation. While the normative syntax to exchange RDF, RDF/XML [13], is an XML dialect itself, there are various other serialisation formats for RDF, such as RDFa [1], a format that allows one to embed RDF within (X)HTML, or non-XML representations such as the more readable Turtle [14] syntax; likewise RDF stores (e.g. YARS2 [54]) normally use their own, proprietary internal representations of triples, that do not relate to XML at all.

**RDF Schema and the Web Ontology Language (OWL)** Although RDF itself is essentially schema-less, additional standards such as RDF Schema and OWL facilitate formal descriptions of the relations between the terms used in an RDF graph: i.e., the predicates in an RDF triple which form edges in an RDF graph (properties) and types of subject or object nodes in an RDF graph (classes). Formal descriptions of these properties and classes can be understood as logical theories, also called ontologies, which allow systems to infer new connections in an RDF graph, or link otherwise unconnected RDF graphs. Standard languages to describe ontologies on the Web are

- RDF Schema [20] – a lightweight ontology language that allows one to describe essentially simple class hierarchies, as well as the domains and ranges of properties; and
- the Web Ontology language (OWL) [108] which was first published in 2004 and recently has been extended with additional useful features in the OWL2 [56] standard.

OWL offers richer means than RDF Schema to define formal relations between classes and properties, such as intersection and union of classes, value restrictions or cardinality restrictions. OWL2 offers even more features such as, for instance, the ability to define keys, property chains, or meta-modeling (i.e., speaking about classes as instances).

**The Rule Interchange Format (RIF)** Although ontology languages such as OWL(2) offer a rich set of constructs to describe relations between RDF terms, these languages are still insufficient to express complex mappings between ontologies, which may better be described in terms of rule languages. The lack of standards in this area had been addressed by several proposals for rule languages on top of RDF, such as the Semantic Web Rule language (SWRL) [62], WRL [6], or N3 [12,15]. These languages offer, for example, support for non-monotonic negation, or rich sets of built-in functions. The importance of rule languages – also outside the narrow use case of RDF rules – has finally lead to the establishment of another W3C working group in 2005 to standardise a generic Rule Interchange Format (RIF). RIF has recently reached proposed recommendation status and will soon be a W3C recommendation. The standard comprises several dialects such as (i)

RIF Core [17], a minimal dialect close to Datalog, (ii) the RIF Basic Logic Dialect (RIF-BLD) [18] which offers the expressive features of Horn rules, and also (iii) a production rules dialect (RIF-PRD) [35]. A set of standard datatypes as well as built-in functions and predicates (RIF-DTB) are defined in a separate document [92]. The relation of RIF to OWL and RDF is detailed in another document [31] that defines the formal semantics of combinations of RIF rule sets with RDF graphs and OWL ontologies.

**Query and Transformation Language: SPARQL** Finally, a crucial puzzle piece which pushed the recent wide uptake of Semantic Web technologies at large was the availability of a standard query language for RDF, namely SPARQL [97], which plays the same role for the Semantic Web as SQL does for relational data. SPARQL's syntax is roughly inspired by Turtle [14] and SQL [109], providing basic means to query RDF such as unions of conjunctive queries, value filtering, optional query parts, as well as slicing and sorting results. The recently re-chartered SPARQL1.1 W3C working group[1] aims at extending the original SPARQL language by commonly requested features such as aggregates, sub-queries, negation, and path expressions.

## 2 Scientific foundations for Semantic Web Standards

The work in the respective standardisation groups is partially still ongoing or only finished very recently. In parallel, there has been plenty of work in the scientific community to define the formal underpinnings for these standards:

- The logical foundations and properties of RDF and RDF Schema have been investigated in detail [83,52,89]. Correspondence of the formal semantics of RDF and RDF Schema [55] with Datalog and First-order logic have been studied in the literature [21,22,66].
- The semantics of standard fragments of OWL have been defined in terms of expressive Description Logics such as $\mathcal{SHOIN}$(D) (OWL DL) [61] or $\mathcal{SROIQ}$(D) (OWL2DL) [60], and the research on OWL has significantly influenced the Description Logics community over the past years: for example, in defining tractable fragments like the $\mathcal{EL}$ [8,9] family of Description Logics, or fragments that allow for reducing basic reasoning tasks to query answering in SQL, such as the DL-Lite family of Description Logics [26]. Other fragments of OWL and OWL2 have been defined in terms of Horn rules such as DLP [51], OWL$^{-}$ [34], pD* [110], or Horn-SHIQ [72]. In fact, the new OWL2 specification defines tractable fragments of OWL based on these results: namely, OWL2EL, OWL2QL, and OWL2RL [79].
- The semantics of RIF builds on foundations such as Frame Logic [70] and Datalog. RIF borrows, e.g., notions of Datalog safety from the scientific literature to define fragments with finite minimal models despite the presence of built-ins: the *strongly-safe* fragment of RIF Core [17, Section 6.2] is inspired by a similar safety condition defined by Eiter, Schindlauer, et al. [39,103]. In fact, the closely related area of decidable subsets of Datalog and answer set programs with function symbols is a very active field of research [10,42,25].
- The formal semantics of SPARQL is also very much inspired by academic results, such as by the seminal papers of Pérez et al. [85,86]. Their work further lead to refined results on equivalences within SPARQL [104] and on the relation of SPARQL to Datalog [91,90]. Angles and Gutierrez [7] later showed that SPARQL has exactly the expressive power of non-recursive safe Datalog with negation.

---

[1] `http://www.w3.org/2009/sparql/wiki`

Likewise, the scientific community has identified and addressed gaps between the Semantic Web standards and the formal paradigms they are based on, which we want turn to next.

## 3 Gaps in the Semantic Web Architecture

Although the standards that make up the Semantic Web architecture have all been established by the W3C, they do not always integrate smoothly, indeed these standards had yet to prove useful "in the wild", i.e., to be applied on real Web data. Particularly, the following significant gaps have been identified in various works over the past years:

**Gap 1: XML vs. RDF** The jump from XML, which is a mere syntax format, to RDF, which is more declartive in nature, is not trivial, but needs to be addressed by appropriate – yet missing – transformation languages for exchanging information between RDF-based and XML-based applications.

**Gap 2: RDF vs. OWL** The clean conceptual model of Description Logics underlying the OWL semantics is not necessarily applicable directly to all RDF data, particularly to messy, potentially inconsistent data as found on the Web.

**Gap 3: RDF/OWL vs. Rules/RIF** There are several theoretical and practical concerns in combining ontologies and rules, such as decidability issues or how to merge classical open world reasoning with non-monotonic closed world inference. The current RIF specification leaves many of these questions open, subject to ongoing research.

**Gap 4: SPARQL vs. RDF Schema/RIF/OWL** Query answering over ontologies and rules and subtopics such as the semantics of SPARQL queries over RDF Schema and OWL ontologies, or querying over combinations of ontologies with RIF rulesets are still neglected by the current standards.

In the following, we will discuss these gaps in more depth, point out how they have been addressed in scientific works so far, including own contributions.

**Gap 1: XML vs. RDF** Although RDF's original normative syntax is an XML dialect, it proves impractical to view an RDF graph as an XML document: e.g., when trying to transform XML data in a custom format into RDF (lifting) or, respectively, RDF data into a specific XML schema (lowering) as may be required by a Web service: while W3C's SAWSDL [44] an GRDDL [29] working groups originally proposed XSLT for these tasks, the various ambiguous formats that RDF/XML can take to represent the same graph form an obstacle for defining uniform transformations [3]: to some extent, treating an RDF graph as an XML document contradicts the declarative nature of RDF. Several proposals to overcome the limitations in lifting and lowering by XSLT include (i) compiling SPARQL queries into XSLT [50], (ii) sequential applications of SPARQL and XSLT queries (via the intermediate step of SPARQL's result format [28], another XML format), or (iii) the extension of XSLT by special RDF access features [114] or SPARQL blocks [16]. Our own proposal – XSPARQL [3,2] – is a new language integrating SPARQL and XQuery; this approach has the advantage of blending two languages that are conceptually very similar and facilitates more concise translations than the previous approaches. XSPARQL has recently been acknowledged as a member submission by the W3C [95,71,75,84].

**Gap 2: RDF vs. OWL** There is a certain "schism" between the core Semantic Web and Description Logics communities on what OWL shall be: the description of an ontology in RDF for RDF data, or an RDF exchange format for Description Logic theories. This schism manifests itself in the W3C's two orthogonal semantic specifications for OWL: OWL2's RDF-based semantics [105], which directly builds upon RDF's model-theoretic semantics [55], and OWL2's direct semantics [80], which builds upon the Description Logics $\mathcal{SROIQ}$ but is not defined for all RDF graphs. Both of them address different use cases; however, particular analyses on Web Data have shown [11,58] that pure OWL(2) in its Description Logics based semantics is not practically applicable: (i) in published Web data we find a lot of non-DL ontologies [11], which only leave to apply the RDF-based semantics; (ii) data and ontologies found on the Web spread across different sources contain a lot of inconsistencies, which – in case one aims to still make sense out of this data – prohibits complete reasoning using Description Logics [58]; (iii) finally, current DL reasoners cannot deal with the amounts of instance data found on the Web, which is in the order of billions of statements. Our own most recent approach – SAOR (Scalable Authoritative OWL Reasoner) [59] – aims at addressing these problems. SAOR provides incomplete, but arguably meaningful inferences over huge data sets crawled from the Web, based on rule-based OWL reasoning inspired by earlier approaches such as pD*[110], with further cautious modifications. Hogan and Decker [57] have later compared this approach to the new standard rule-based OWL2RL [79] profile, coming to the conclusion that OWL2RL, as a maximal fragment of OWL2 that can be formalised purely with Horn rules, runs into similar problems as Description Logics reasoning when taken as a basis for reasoning over Web data without the further modifications proposed in SAOR. An orthogonal approach to reason with real Web data [36] – also proposed by the author of this work together with Delbru, Tummarello and Decker – is likewise based on pD*, but applies inference in a modular fashion per dataset rather than over entire Web crawls.

**Gap 3: RDF/OWL vs. Rules/RIF** Issues on combining RDF and/or OWL with rules, and particularly with rule sets expressed in RIF, have so far mostly been discussed on a theoretical level, perhaps because there has not yet been time enough for meaningful adoption of RIF on the Web.

One strand of these discussions is concerned with extending RDF with rules and constraints, in terms of either suggesting new non-standard rule languages for RDF to publish such rules [106,15,5,6,4], or theoretical considerations such as redundancy elimination with rules and constraints on top of RDF [78,88]. An interesting side issue here concerns rule languages that allow existentials in the head such as RDFLog [23], or more recently Datalog$^{+/-}$ [24], which may in fact be viewed as a viable alternative or complement to purely Description Logics based ontology languages. Non-monotonicity – which is not considered in OWL, but is available in most of the suggested rule languages for RDF [5,15,6] by incorporating a form of "negation as failure" – has sparked a lot of discussions in the Semantic Web community, since it was viewed as inadequate for an open environment such as the Web by some, whereas others (including the author of the present work) argued that "scoped negation" [69,93] – that is, non-monotonic negation applied over a fixed, scoped part of the Web – was very useful for many Web data applications. This is closely related to what Etzioni et al. [43] called the "local closed world assumption" in earlier work.

Another quite significant strand of research has developed on the theoretical combination of Description Logics and (non-monotonic) rules in a joint logical framework. While the naïve combination of even Horn rules without function symbols and ontologies in quite inexpressive Description Logics loses the desirable decidability properties of the latter [74], there have been several proposals for decidable fragments of this combination [51,82,72] or even extending the idea of such decidable combinations to rules with non-monotonic negation [98,99,101,81,77]. Another decidable approach was to define the semantic interplay between ontologies and rules via a narrow, query-like interface within rule bodies [40]. Aside from considerations about decidability, there have been several proposals for what would be the right logical framework to embed combinations of classical logical theories (which DL ontologies fall into) and non-monotonic rule languages. These include approaches based on MKNF [81], FO-AEL [32], or Quantified Equilibrium Logics (QEL) [33]. For an overview of issues concerned with combining ontologies and rules, we also refer to surveys of existing approaches in [38,37,100].

As a side note, it should be mentioned that rule-based/resolution-based reasoning has been very successfully applied in implementing Description Logics or OWL reasoners in approaches such as KAON2 [63] and DLog [76] which significantly outperform tableaux-based DL reasoners on certain problems (particularly instance reasoning).

**Gap 4: SPARQL vs. RDF Schema/RIF/OWL** SPARQL has in its official specification only been defined as a query language over RDF graphs, not taking into account RDF Schema, OWL ontologies or RIF rule sets. Although the official specification defines frame conditions for extending SPARQL by higher entailment regimes [97, Section 12.6], few works have actually instantiated this mechanism and defined how SPARQL should handle ontologies and rule sets.

As for OWL, conjunctive query answering over expressive description logics is a topic of active research in the Description Logics Community, with important insights only being very recent [41,47,46,73], none of which yet having covered the Description Logics underlying OWL(2), $\mathcal{SHOIN}$(D) and $\mathcal{SROIQ}$(D). Answering full SPARQL queries on top of OWL has only preliminarily been addressed in the scientific community [107,67] so far.

In terms of SPARQL on top of RDF in combination with rule sets, the choices are more obvious. Firstly, as mentioned above, SPARQL itself can be translated to non-recursive rules – more precisely into non-recursive Datalog with negation [91,7]. Secondly, expanding on the translation from [91], additional RDF rule sets that guarantee a finite closure, such as Datalog style rules on top of RDF, can be allowed, covering a significant subset of RIF or rule-based approximations of RDFS and OWL [65,64].

One should mention here that certain SPARQL queries themselves may be read as rules: that is, SPARQL's CONSTRUCT queries facilitate the generation of new RDF triples (defined in a CONSTRUCT template that plays the role of the rule head), based on the answers to a graph pattern (that plays the role of a rule body). This idea has been the basis for proposals to extend RDF to so-called Networked Graphs [102] or Extended RDF graphs [96], that enable the inclusion of implicit knowledge defined as SPARQL CONSTRUCT queries. We have also proposed to extend RDF graphs in such fashions as an expressive means to define ontology mappings [96].

The recently started W3C SPARQL1.1 working group has published a working draft summarising first results on defining an OWL entailment regime for SPARQL [49], which, although worth to be mentioned, will not necessarily encompass full conjunctive queries with non-distinguished variables.

## 4 Conclusions

The present paper tried to summarise current developments and trends in terms of Semantic Web standards, highlighting gaps between these standards and surveying scientific works that have provided foundations to these standards or promise to close these gaps. We hope this subjective selection serves as an entry point for the interested reader. The work presented has been supported in parts by (i) Science Foundation Ireland – under the Líon (SFI/02/CE1/I131) and Líon-2 (SFI/08/CE/I1380) projects. The author especially thanks all co-authors of own works cited in this paper.

# References

1. Ben Adida, Mark Birbeck, Shane McCarron, and Steven Pemberton. RDFa in XHTML: Syntax and Processing. W3C recommendation, W3C, October 2008. Available at `http://www.w3.org/TR/rdfa-syntax/`.
2. Waseem Akhtar, Jacek Kopecký, Thomas Krennwallner, and Axel Polleres. XSPARQL: Traveling between the XML and RDF worlds – and avoiding the XSLT pilgrimage. Technical Report DERI-TR-2007-12-14, DERI Galway, 2007. Available at `http://www.deri.ie/fileadmin/documents/TRs/DERI-TR-2007-12-14.pdf`.
3. Waseem Akhtar, Jacek Kopecky, Thomas Krennwallner, and Axel Polleres. XSPARQL: Traveling between the XML and RDF worlds – and avoiding the XSLT pilgrimage. In *Proceedings of the 5th European Semantic Web Conference (ESWC2008)*, pages 432–447, Tenerife, Spain, June 2008. Springer.
4. Anastasia Analyti, Grigoris Antoniou, and Carlos Viegas Damásio. A principled framework for modular web rule bases and its semantics. In *Proceedings of the 11th International Conference on Principles of Knowledge Representation and Reasoning (KR'08)*, pages 390–400, 2008.
5. Anastasia Analyti, Grigoris Antoniou, Carlos Viegas Damasio, and Gerd Wagner. Extended RDF as a semantic foundation of rule markup languages. *Journal of Artificial Intelligence Research*, 32:37–94, 2008.
6. Jürgen Angele, Harold Boley, Jos de Bruijn, Dieter Fensel, Pascal Hitzler, Michael Kifer, Reto Krummenacher, Holger Lausen, Axel Polleres, and Rudi Studer. Web Rule Language (WRL), September 2005. W3C member submission.
7. Renzo Angles and Claudio Gutierrez. The expressive power of sparql. In *International Semantic Web Conference (ISWC 2008)*, volume 5318 of *Lecture Notes in Computer Science*, pages 114–129, Karlsruhe, Germany, 2008. Springer.
8. Franz Baader. Terminological cycles in a description logic with existential restrictions. In *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI2003)*, pages 325–330, Acapulco, Mexico, August 2003.
9. Franz Baader, Sebastian Brandt, and Carsten Lutz. Pushing the el envelope. In *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence (IJCAI2005)*, pages 364–369, Edinburgh, Scotland, UK, July 2005. Professional Book Center.
10. Sabrina Baselice, Piero A. Bonatti, and Giovanni Criscuolo. On finitely recursive programs. *TPLP*, 9(2):213–238, 2009.
11. Sean Bechhofer and Raphael Volz. Patching syntax in OWL ontologies. In *International Semantic Web Conference (ISWC 2004)*, pages 668–682, Hiroshima, Japan, November 2004.
12. Dave Beckett and Tim Berners-Lee. Turtle – Terse RDF Triple Language. W3C team submission, W3C, January 2008. Available at `http://www.w3.org/TeamSubmission/turtle/`.
13. Dave Beckett and Brian McBride. RDF/XML Syntax Specification (Revised). W3C recommendation, W3C, February 2004. Available at `http://www.w3.org/TR/REC-rdf-syntax/`.
14. Tim Berners-Lee and Dan Connolly. Notation3 (N3): A readable RDF syntax. W3C team submission, W3C, January 2008. Available at `http://www.w3.org/TeamSubmission/n3/`.
15. Tim Berners-Lee, Dan Connolly, Lalana Kagal, Yosi Scharf, and Jim Hendler. N3logic: A logical framework for the world wide web. *Theory and Practice of Logic Programming*, 8(3):249–269, 2008.
16. Diego Berrueta, Jose E. Labra, and Ivan Herman. XSLT+SPARQL : Scripting the Semantic Web with SPARQL embedded into XSLT stylesheets. In Chris Bizer, Sören Auer, Gunnar Aastrand Grimmes, and Tom Heath, editors, *4th Workshop on Scripting for the Semantic Web*, Tenerife, June 2008.
17. Harold Boley, Gary Hallmark, Michael Kifer, Adrian Paschke, Axel Polleres, and Dave Reynolds. RIF Core Dialect. W3C proposed recommendation, W3C, May 2010. Available at `http://www.w3.org/TR/2010/PR-rif-core-20100511/`.
18. Harold Boley and Michael Kifer. RIF Basic Logic Dialect. W3C proposed recommendation, W3C, May 2010. Available at `http://www.w3.org/TR/2010/PR-rif-bld-20100511/`.
19. Tim Bray, Jean Paoli, and C.M. Sperberg-McQueen. Extensible Markup Language (XML) 1.0. W3C Recommendation, W3C, February 1998. Available at `http://www.w3.org/TR/1998/REC-xml-19980210`.
20. Dan Brickley, R. Guha, and Brian McBride (eds.). RDF Vocabulary Description Language 1.0: RDF Schema. Technical report, W3C, February 2004. W3C Recommendation.
21. Jos de Bruijn, Enrico Franconi, and Sergio Tessaris. Logical reconstruction of normative RDF. In *OWL: Experiences and Directions Workshop (OWLED-2005)*, Galway, Ireland, November 2005.
22. Jos de Bruijn and Stijn Heymans. Logical foundations of (e)RDF(S): Complexity and reasoning. In *Proceedings of the 6th International Semantic Web Conference and 2nd Asian Semantic Web Conference (ISWC2007+ASWC2007)*, number 4825 in Lecture Notes in Computer Science, pages 86–99, Busan, Korea, November 2007. Springer.
23. François Bry, Tim Furche, Clemens Ley, Benedikt Linse, and Bruno Marnette. RDFLog: It's like datalog for RDF. In *Proceedings of 22nd Workshop on (Constraint) Logic Programming, Dresden (30th September–1st October 2008)*, 2008.
24. Andrea Calì, Georg Gottlob, and Thomas Lukasiewicz. Tractable query answering over ontologies with datalog$^{+/-}$. In *Proceedings of the 22nd International Workshop on Description Logics (DL 2009)*, Oxford, UK, July 2009.
25. Francesco Calimeri, Susanna Cozza, Giovambattista Ianni, and Nicola Leone. Magic sets for the bottom-up evaluation of finitely recursive programs. In Esra Erdem, Fangzhen Lin, and Torsten Schaub, editors, *Logic Programming and Nonmonotonic Reasoning, 10th International Conference (LPNMR 2009)*, volume 5753 of *Lecture Notes in Computer Science*, pages 71–86, Potsdam, Germany, September 2009. Springer.
26. Diego Calvanese, Giuseppe De Giacomo, Domenico Lembo, Maurizio Lenzerini, and Riccardo Rosati. Tractable reasoning and efficient query answering in description logics: The *dl-lite* family. *Journal of Automated Reasoning*, 39(3):385–429, 2007.
27. Don Chamberlin, Jonathan Robie, Scott Boag, Mary F. Fernández, Jérôme Siméon, and Daniela Florescu. XQuery 1.0: An XML Query Language. W3C recommendation, W3C, January 2007. W3C Recommendation, available at `http://www.w3.org/TR/xquery/`.
28. Kendall Grant Clark, Lee Feigenbaum, and Elias Torres. SPARQL Protocol for RDF. W3C recommendation, W3C, January 2008. Available at `http://www.w3.org/TR/rdf-sparql-protocol/`.
29. Dan Connolly. Gleaning Resource Descriptions from Dialects of Languages (GRDDL). W3C Recommendation, W3C, September 2007. Available at `http://www.w3.org/TR/sawsdl/`.

30. Stéphane Corlosquet, Renaud Delbru, Tim Clark, Axel Polleres, and Stefan Decker. Produce and consume linked data with Drupal! In Abraham Bernstein, David R. Karger, Tom Heath, Lee Feigenbaum, Diana Maynard, Enrico Motta, and Krishnaprasad Thirunarayan, editors, *Proceedings of the 8th International Semantic Web Conference (ISWC 2009)*, volume 5823 of *Lecture Notes in Computer Science*, pages 763–778, Washington DC, USA, October 2009. Springer.
31. Jos de Bruijn. RIF RDF and OWL Compatibility. W3C proposed recommendation, W3C, May 2010. Available at `http://www.w3.org/TR/2010/PR-rif-rdf-owl-20100511/`.
32. Jos de Bruijn, Thomas Eiter, Axel Polleres, and Hans Tompits. Embedding non-ground logic programs into autoepistemic logic for knowledge-base combination. In *Twentieth International Joint Conference on Artificial Intelligence (IJCAI'07)*, pages 304–309, Hyderabad, India, January 2007. AAAI.
33. Jos de Bruijn, David Pearce, Axel Polleres, and Agustín Valverde. A semantical framework for hybrid knowledge bases. *Knowledge and Information Systems*, Special Issue: RR 2007, 2010. Accepted for publication.
34. Jos de Bruijn, Axel Polleres, Rubén Lara, and Dieter Fensel. OWL$^-$. Final draft d20.1v0.2, WSML, 2005.
35. Christian de Sainte Marie, Gary Hallmark, and Adrian Paschke. RIF Production Rule Dialect. W3C proposed recommendation, W3C, May 2010. Available at `http://www.w3.org/TR/2010/PR-rif-prd-20100511/`.
36. Renaud Delbru, Axel Polleres, Giovanni Tummarello, and Stefan Decker. Context dependent reasoning for semantic documents in sindice. In *Proceedings of the 4th International Workshop on Scalable Semantic Web Knowledge Base Systems (SSWS 2008)*, Karlsruhe, Germany, October 2008.
37. Thomas Eiter, Giovambattista Ianni, Thomas Krennwallner, and Axel Polleres. Rules and ontologies for the semantic web. In Cristina Baroglio, Piero A. Bonatti, Jan Maluszynski, Massimo Marchiori, Axel Polleres, and Sebastian Schaffert, editors, *Reasoning Web 2008*, volume 5224 of *Lecture Notes in Computer Science*, pages 1–53. Springer, San Servolo Island, Venice, Italy, September 2008.
38. Thomas Eiter, Giovambattista Ianni, Axel Polleres, Roman Schindlauer, and Hans Tompits. Reasoning with rules and ontologies. In P. Barahona et al., editor, *Reasoning Web 2006*, volume 4126 of *Lecture Notes in Computer Science*, pages 93–127. Springer, September 2006.
39. Thomas Eiter, Giovambattista Ianni, Roman Schindlauer, and Hans Tompits. Effective integration of declarative rules with external evaluations for semantic-web reasoning. In *Proceedings of the 3rd European Semantic Web Conference (ESWC2006)*, volume 4011 of *LNCS*, pages 273–287, Budva, Montenegro, June 2006. Springer.
40. Thomas Eiter, Thomas Lukasiewicz, Roman Schindlauer, and Hans Tompits. Combining answer set programming with description logics for the semantic web. In *Proceedings of the Ninth International Conference on Principles of Knowledge Representation and Reasoning (KR'04)*, Whistler, Canada, 2004. AAAI Press.
41. Thomas Eiter, Carsten Lutz, Magdalena Ortiz, and Mantas Simkus. Query answering in description logics with transitive roles. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI 2009)*, pages 759–764, Pasadena, California, USA, July 2009.
42. Thomas Eiter and Mantas Simkus. FDNC: Decidable nonmonotonic disjunctive logic programs with function symbols. *ACM Trans. Comput. Log.*, 11(2), 2010.
43. Oren Etzioni, Keith Golden, and Daniel Weld. Tractable closed world reasoning with updates. In *KR'94: Principles of Knowledge Representation and Reasoning*, pages 178–189, San Francisco, California, 1994. Morgan Kaufmann.
44. Joel Farrell and Holger Lausen. Semantic Annotations for WSDL and XML Schema. W3C Recommendation, W3C, August 2007. Available at `http://www.w3.org/TR/sawsdl/`.
45. Dieter Fensel, Holger Lausen, Axel Polleres, Jos de Bruijn, Michael Stollberg, Dumitru Roman, and John Domingue. *Enabling Semantic Web Services : The Web Service Modeling Ontology*. Springer, 2006.
46. Birte Glimm, Ian Horrocks, and Ulrike Sattler. Unions of conjunctive queries in SHOQ. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Eleventh International Conference, KR 2008*, pages 252–262, Sydney, Australia, September 2008. AAAI Press.
47. Birte Glimm, Carsten Lutz, Ian Horrocks, and Ulrike Sattler. Conjunctive query answering for the description logic SHIQ. *J. Artif. Intell. Res. (JAIR)*, 31:157–204, 2008.
48. Status QIO: Conjunctive Query Entailment is Decidable. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Twelfth International Conference, KR 2010*, pages 225–235, Toronto, Canada, May 2010. AAAI Press.
49. Birte Glimm, Chimezie Ogbuji, Sandro Hawke, Ivan Herman, Bijan Parsia, Axel Polleres, and Andy Seaborne. SPARQL 1.1 Entailment Regimes. W3C working draft, W3C, May 2010. Available at `http://www.w3.org/TR/sparql11-entailment/`.
50. Sven Groppe, Jinghua Groppe, Volker Linnemann, Dirk Kukulenz, Nils Hoeller, and Christoph Reinke. Embedding SPARQL into XQuery/XSLT. In *Proceedings of the 2008 ACM Symposium on Applied Computing (SAC)*, pages 2271–2278, Fortaleza, Ceara, Brazil, March 2008. ACM.
51. Benjamin N. Grosof, Ian Horrocks, Raphael Volz, and Stefan Decker. Description logic programs: Combining logic programs with description logic. In *12th International Conference on World Wide Web (WWW'03)*, pages 48–57, Budapest, Hungary, 2003. ACM.
52. Claudio Gutiérrez, Carlos A. Hurtado, and Alberto O. Mendelzon. Foundations of Semantic Web Databases. In *Proceedings of the Twenty-third ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS 2004)*, pages 95–106, Paris, France, 2004. ACM.
53. Andreas Harth, Katja Hose, Marcel Karnstedt, Axel Polleres, Kai-Uwe Sattler, and Jürgen Umbrich. Data summaries for on-demand queries over linked data. In *Proceedings of the 19th World Wide Web Conference (WWW2010)*, Raleigh, NC, USA, April 2010. ACM Press. Technical report version available at `http://www.deri.ie/fileadmin/documents/DERI-TR-2009-11-17.pdf`.
54. Andreas Harth, Jürgen Umbrich, Aidan Hogan, and Stefan Decker. YARS2: A federated repository for querying graph structured data from the web. In *6th International Semantic Web Conference, 2nd Asian Semantic Web Conference*, pages 211–224, 2007.
55. Patrick Hayes. RDF semantics. Technical report, W3C, February 2004. W3C Recommendation.
56. Pascal Hitzler, Markus Krötzsch, Bijan Parsia, Peter F. Patel-Schneider, and Sebastian Rudolph. OWL 2 Web Ontology Language Primer. W3C recommendation, W3C, October 2009. Available at `http://www.w3.org/TR/owl2-primer/`.

57. Aidan Hogan and Stefan Decker. On the ostensibly silent 'W' in OWL 2 RL. In *Web Reasoning and Rule Systems – Third International Conference, RR 2009*, pages 118–134, 2009.
58. Aidan Hogan, Andreas Harth, Alexandre Passant, Stefan Decker, and Axel Polleres. Weaving the pedantic web. In *3rd International Workshop on Linked Data on the Web (LDOW2010) at WWW2010*, Raleigh, USA, April 2010.
59. Aidan Hogan, Andreas Harth, and Axel Polleres. Scalable authoritative OWL reasoning for the Web. *International Journal on Semantic Web and Information Systems*, 5(2):49–90, 2009.
60. Ian Horrocks, Oliver Kutz, and Ulrike Sattler. The even more irresistible SROIQ. In *Proceedings of the Tenth International Conference on Principles of Knowledge Representation and Reasoning (KR'06)*, pages 57–67. AAAI Press, 2006.
61. Ian Horrocks and Peter F. Patel-Schneider. Reducing OWL entailment to description logic satisfiability. *Journal of Web Semantics*, 1(4):345–357, 2004.
62. Ian Horrocks, Peter F. Patel-Schneider, Harold Boley, Said Tabet, Benjamin Grosof, and Mike Dean. SWRL: A Semantic Web Rule Language Combining OWL and RuleML, May 2004. W3C member submission.
63. Ullrich Hustadt, Boris Motik, and Ulrike Sattler. Reducing shiq-description logic to disjunctive datalog programs. In *Proceedings of the Ninth International Conference on Principles of Knowledge Representation and Reasoning (KR'04)*, pages 152–162, Whistler, Canada, 2004. AAAI Press.
64. Giovambattista Ianni, Thomas Krennwallner, Alessandra Martello, and Axel Polleres. Dynamic querying of mass-storage RDF data with rule-based entailment regimes. In Abraham Bernstein, David R. Karger, Tom Heath, Lee Feigenbaum, Diana Maynard, Enrico Motta, and Krishnaprasad Thirunarayan, editors, *International Semantic Web Conference (ISWC 2009)*, volume 5823 of *Lecture Notes in Computer Science*, pages 310–327, Washington DC, USA, October 2009. Springer.
65. Giovambattista Ianni, Thomas Krennwallner, Alessandra Martello, and Axel Polleres. A rule system for querying persistent RDFS data. In *Proceedings of the 6th European Semantic Web Conference (ESWC2009)*, Heraklion, Greece, May 2009. Springer. Demo Paper.
66. Giovambattista Ianni, Alessandra Martello, Claudio Panetta, and Giorgio Terracina. Efficiently querying RDF(S) ontologies with Answer Set Programming. *Journal of Logic and Computation (Special issue)*, 19(4):671–695, August 2009.
67. Yixin Jing, Dongwon Jeong, and Doo-Kwon Baik. SPARQL graph pattern rewriting for OWL-DL inference queries. *Knowl. Inf. Syst.*, 20(2):243–262, 2009.
68. Michael Kay. XSL Transformations (XSLT) Version 2.0 . W3C Recommendation, W3C, January 2007. Available at `http://www.w3.org/TR/xslt`20.
69. Michael Kifer. Nonmonotonic reasoning in FLORA-2. In *8th Int'l Conf. on Logic Programming and Nonmonotonic Reasoning (LPNMR'05)*, Diamante, Italy, 2005. Invited Paper.
70. Michael Kifer, Georg Lausen, and James Wu. Logical foundations of object-oriented and frame-based languages. *Journal of the ACM*, 42(4):741–843, 1995.
71. Thomas Krennwallner, Nuno Lopes, and Axel Polleres. XSPARQL: Semantics, January 2009. W3C member submission.
72. Markus Krötzsch, Sebastian Rudolph, and Pascal Hitzler. Complexity boundaries for horn description logics. In *Proceedings of the Twenty-Second AAAI Conference on Artificial Intelligence (AAAI)*, pages 452–457, Vancouver, British Columbia, Canada, July 2007.
73. Markus Krötzsch, Sebastian Rudolph, and Pascal Hitzler. Conjunctive queries for a tractable fragment of OWL 1.1. In *Proceedings of the 6th International Semantic Web Conference and 2nd Asian Semantic Web Conference, ISWC 2007 + ASWC 2007*, pages 310–323, Busan, Korea, November 2007.
74. Alon Y. Levy and Marie-Christine Rousset. Combining horn rules and description logics in CARIN. *Artificial Intelligence*, 104:165–209, 1998.
75. Nuno Lopes, Thomas Krennwallner, Axel Polleres, Waseem Akhtar, and Stéphane Corlosquet. XSPARQL: Implementation and Test-cases, January 2009. W3C member submission.
76. Gergely Lukácsy and Péter Szeredi. Efficient description logic reasoning in Prolog: the DLog system. *Theory and Practice of Logic Programming*, 9(3):343–414, 2009.
77. Thomas Lukasiewicz. A novel combination of answer set programming with description logics for the semantic web. *IEEE Transactions on Knowledge and Data Engineering (TKDE)*, 2010. In press.
78. Michael Meier. Towards Rule-Based Minimization of RDF Graphs under Constraints. In *Proc. RR'08*, volume 5341 of *LNCS*, pages 89–103. Springer, 2008.
79. Boris Motik, Bernardo Cuenca Grau, Ian Horrocks, Zhe Wu, Achille Fokoue, Casrsten Lutz, Diego Calvanese, Jeremy Carroll, Guiseppe De Giacomo, Jim Hendler, Ivan Herman, Bijan Parsia, Peter F. Patel-Schneider, Alan Ruttenberg, Uli Sattler, and Michael Schneider. OWL 2 Web Ontology Language Profiles. W3C recommendation, W3C, October 2009. Available at `http://www.w3.org/TR/owl2-profiles/`.
80. Boris Motik, Peter F. Patel-Schneider, Bernardo Cuenca Grau, Ian Horrocks, Bijan Parsia, and Uli Sattler. OWL 2 Web Ontology Language Direct Semantics. W3C recommendation, W3C, October 2009. Available at `http://www.w3.org/TR/owl2-direct-semantics/`.
81. Boris Motik and Riccardo Rosati. A faithful integration of description logics with logic programming. In *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI-07)*, pages 477–482, Hyderabad, India, January 6–12 2007. AAAI.
82. Boris Motik, Ulrike Sattler, and Rudi Studer. Query answering for OWL-DL with rules. *Journal of Web Semantics*, 3(1):41–60, 2005.
83. Sergio Muñoz, Jorge Pérez, and Claudio Gutiérrez. Minimal deductive systems for RDF. In Enrico Franconi, Michael Kifer, and Wolfgang May, editors, *Proceedings of the 4th European Semantic Web Conference (ESWC2007)*, volume 4519 of *Lecture Notes in Computer Science*, pages 53–67, Innsbruck, Austria, June 2007. Springer.
84. Alexandre Passant, Jacek Kopecký, Stéphane Corlosquet, Diego Berrueta, Davide Palmisano, and Axel Polleres. XSPARQL: Use cases, January 2009. W3C member submission.
85. Jorge Pérez, Marcelo Arenas, and Claudio Gutierrez. Semantics and complexity of sparql. In *International Semantic Web Conference (ISWC 2006)*, pages 30–43, 2006.

86. Jorge Pérez, Marcelo Arenas, and Claudio Gutierrez. Semantics and complexity of sparql. *ACM Transactions on Database Systems*, 34(3):Article 16 (45 pages), 2009.
87. Danh Le Phuoc, Axel Polleres, Giovanni Tummarello, Christian Morbidoni, and Manfred Hauswirth. Rapid semantic web mashup development through semantic web pipes. In *Proceedings of the 18th World Wide Web Conference (WWW2009)*, pages 581–590, Madrid, Spain, April 2009. ACM Press.
88. Reinhard Pichler, Axel Polleres, Sebastian Skritek, and Stefan Woltran. Minimising RDF graphs under rules and constraints revisited. In *4th Alberto Mendelzon Workshop on Foundations of Data Management*, May 2010. To appear, technical report version available at `http://www.deri.ie/fileadmin/documents/DERI-TR-2010-04-23.pdf`.
89. Reinhard Pichler, Axel Polleres, Fang Wei, and Stefan Woltran. Entailment for domain-restricted RDF. In *Proceedings of the 5th European Semantic Web Conference (ESWC2008)*, pages 200–214, Tenerife, Spain, June 2008. Springer.
90. Axel Polleres. SPARQL Rules! Technical Report GIA-TR-2006-11-28, Universidad Rey Juan Carlos, Móstoles, Spain, 2006. Available at `http://www.polleres.net/TRs/GIA-TR-2006-11-28.pdf`.
91. Axel Polleres. From SPARQL to rules (and back). In *Proceedings of the 16th World Wide Web Conference (WWW2007)*, pages 787–796, Banff, Canada, May 2007. ACM Press. Extended technical report version available at `http://www.polleres.net/TRs/GIA-TR-2006-11-28.pdf`, slides available at `http://www.polleres.net/publications/poll-2007www-slides.pdf`.
92. Axel Polleres, Harold Boley, and Michael Kifer. RIF Datatypes and Built-Ins 1.0. W3C proposed recommendation, W3C, May 2010. Available at `http://www.w3.org/TR/2010/PR-rif-dtb-20100511/`.
93. Axel Polleres, Cristina Feier, and Andreas Harth. Rules with contextually scoped negation. In *Proceedings of the 3rd European Semantic Web Conference (ESWC2006)*, volume 4011 of *Lecture Notes in Computer Science*, Budva, Montenegro, June 2006. Springer.
94. Axel Polleres and David Huynh, editors. *Journal of Web Semantics, Special Issue: The Web of Data*, volume 7(3). Elsevier, 2009.
95. Axel Polleres, Thomas Krennwallner, Nuno Lopes, Jacek Kopecký, and Stefan Decker. XSPARQL Language Specification, January 2009. W3C member submission.
96. Axel Polleres, François Scharffe, and Roman Schindlauer. SPARQL++ for mapping between RDF vocabularies. In *OTM 2007, Part I : Proceedings of the 6th International Conference on Ontologies, DataBases, and Applications of Semantics (ODBASE 2007)*, volume 4803 of *Lecture Notes in Computer Science*, pages 878–896, Vilamoura, Algarve, Portugal, November 2007. Springer.
97. Eric Prud′hommeaux and Andy Seaborne. SPARQL Query Language for RDF. W3C recommendation, W3C, January 2008. Available at `http://www.w3.org/TR/rdf-sparql-query/`.
98. Riccardo Rosati. On the decidability and complexity of integrating ontologies and rules. *Journal of Web Semantics*, 3(1):61–73, 2005.
99. Riccardo Rosati. Semantic and computational advantages of the safe integration of ontologies and rules. In *Proceedings of the Third International Workshop on Principles and Practice of Semantic Web Reasoning (PPSWR 2005)*, volume 3703 of *Lecture Notes in Computer Science*, pages 50–64. Springer, 2005.
100. Riccardo Rosati. Integrating Ontologies and Rules: Semantic and Computational Issues. In Pedro Barahona, François Bry, Enrico Franconi, Ulrike Sattler, and Nicola Henze, editors, *Reasoning Web, Second International Summer School 2006, Lissabon, Portugal, September 25-29, 2006, Tutorial Lectures*, volume 4126 of *LNCS*, pages 128–151. Springer, September 2006.
101. Riccardo Rosati. $\mathcal{DL}+log$: Tight integration of description logics and disjunctive datalog. In *Proceedings of the Tenth International Conference on Principles of Knowledge Representation and Reasoning (KR'06)*, pages 68–78, 2006.
102. Simon Schenk and Steffen Staab. Networked graphs: A declarative mechanism for SPARQL rules, SPARQL views and RDF data integration on the web. In *Proceedings WWW-2008*, pages 585–594, Beijing, China, 2008. ACM Press.
103. Roman Schindlauer. *Answer-Set Programming for the Semantic Web*. PhD thesis, Vienna University of Technology, December 2006.
104. Michael Schmidt, Michael Meier, and Georg Lausen. Foundations of sparql query optimization. In *13th International Conference on Database Theory (ICDT2010)*, Lausanne, Switzerland, March 2010.
105. Michael Schneider, Jeremy Carroll, Ivan Herman, and Peter F. Patel-Schneider. W3C OWL 2 Web Ontology Language RDF-Based Semantics. W3C recommendation, W3C, October 2009. Available at `http://www.w3.org/TR/owl2-rdf-based-semantics/`.
106. Michael Sintek and Stefan Decker. TRIPLE - A Query, Inference, and Transformation Language for the Semantic Web. In *1st International Semantic Web Conference*, pages 364–378, 2002.
107. Evren Sirin and Bijan Parsia. SPARQL-DL: SPARQL query for OWL-DL. In *Proceedings of the OWLED 2007 Workshop on OWL: Experiences and Directions*, Innsbruck, Austria, June 2007. CEUR-WS.org.
108. Michael K. Smith, Chris Welty, and Deborah L. McGuinness. OWL Web Ontology Language Guide. W3C recommendation, W3C, February 2004. Available at `http://www.w3.org/TR/owl-guide/`.
109. SQL-99. Information Technology - Database Language SQL- Part 3: Call Level Interface (SQL/CLI). Technical Report INCITS/ISO/IEC 9075-3, INCITS/ISO/IEC, October 1999. Standard specification.
110. Herman J. ter Horst. Completeness, decidability and complexity of entailment for RDF Schema and a semantic extension involving the OWL vocabulary. *Journal of Web Semantics*, 3:79–115, 2005.
111. Giorgio Terracina, Nicola Leone, Vincenzino Lio, and Claudio Panetta. Experimenting with recursive queries in database and logic programming systems. *Theory and Practice of Logic Programming*, 8(2):129–165, March 2008.
112. H. S. Thompson, D. Beech, M. Maloney, and N. Mendelsohn. XML Schema Part 1: Structures, 2nd Edition. W3C Recommendation, W3C, October 2004. Available at `http://www.w3.org/TR/xmlschema-1/`.
113. Jacopo Urbani, Spyros Kotoulas, Eyal Oren, and Frank van Harmelen. Scalable distributed reasoning using mapreduce. In *International Semantic Web Conference*, pages 634–649, 2009.
114. Norman Walsh. RDF Twig: Accessing RDF Graphs in XSLT. Presented at Extreme Markup Languages (XML) 2003, Montreal, Canada. Available at `http://rdftwig.sourceforge.net/`.
115. Jesse Weaver and James A. Hendler. Parallel materialization of the finite RDFS closure for hundreds of millions of triples. In *International Semantic Web Conference (ISWC2009)*, pages 682–697, 2009.