

Making Linked Data Reasonable using Description Logics, Part 3

by Mike Bergman - Wednesday, February 18, 2009

<http://www.mkbergman.com/477/making-linked-data-reasonable-using-description-logics-part-3/>



Historical and Research Support for Splitting the TBox and ABox

In [Part 1](#) of this series, I advocated the placement of [linked data](#) in an ABox construct from [description logics \[1\]](#) based on a [separation of concerns](#) argument. In [Part 2](#) of this series, I reinforced that argument from the perspective of the *work* to be done within a knowledge base.

I came to these viewpoints independently. I do not have any special background in these disciplines; I am a recent researcher and practitioner in the field, perhaps akin to a gentleman natural scientist of the 1800s. As these ideas have formed, therefore, I have also attempted to see what some of the noted experts in the field have said and wrote.

Like any other field, there is no common viewpoint or doctrine about these matters. But, there is considerable -- and historic -- support for this viewpoint of splitting the TBox and the ABox for many different reasons. To my knowledge, this viewpoint has yet to be consolidated and applied to linked data. Perhaps this series will help stimulate that discussion.

A Bit of 'Ancient History'

The first specific discussion of this matter I was able to discover (though I suspect it had been discussed in earlier internal forums or papers) was on the W3C's RDF logic mail lists in 2001. While only eight years ago, it does feel a bit like ancient history with regard to the development and understanding of semantic Web languages.

The mail list topic was what role RDF should assume, at a formative point in the language's development. A possibly restricted scope for RDF akin to a relational database or even as an "ABox" was being discussed. (This restriction, of course, was dropped for the more open, free scope of the present RDF.) To help clarify these matters, Jérôme Euzenat [first noted](#) in a thread with Pat Hayes, who would later author

the RDF semantics W3C standards document [\[2\]](#) two years hence, that:

The TBox defines a terminology, i.e. assign terms (which are interpreted as sets) to names, while the ABox asserts formulas generally about individuals. The ABox can use the names assigned in the TBox, while the TBox usually do not refer to those interpreted as individuals in the ABox (but this is evolving). This meets the mathematical delimitation between definitions and assertions (no?!).

Logically this can be cleanly defined by separating the languages. In model theoretic terms, this means that all terms are interpreted as sets (and all assertions can be reduced to inclusion between these sets).

There is an important consequence on that separation (that held at least in first DL languages): the terminology itself cannot be found inconsistent. This is because, it only deals with sets (e.g. the set of things with more than four legs and less than two legs) and the worst that can happen to them is to be empty (but having a whole theory about plenty of empty sets is not inconsistent). On the contrary the ABox can be inconsistent, just because it asserts things (e.g. that there exists something with four legs belonging to the set of things with at most three legs).

This elucidated some discussion about TBox and ABox roles and purposes. [Ian Horrocks](#), one of the lead authors of [DAML+OIL](#) and now [OWL](#), replied [later in that thread](#):

The DL distinction between Tbox and Abox has changed over the years. In the days when reasoning algorithms were less well understood, the split was more or less into that part of the KB [knowledge base] for which (something like) sound and complete (subsumption) reasoning could be performed (the Tbox) and the part for which it could not (the Abox). Nowadays, the distinction is simply the type of axioms: the Tbox consists of (class) concept inclusion axioms (and/or equivalence) - e.g., "C subsumes D" while the Abox consists of individual/tuple membership axioms - e.g., "x is an instance of C" or "<x,y> is an instance of R". For most DLs this is a useful distinction as a simpler algorithm can be used for reasoning with a Tbox alone (no individual names).

. . . Even today, where Tboxes are assumed to contain arbitrary inclusion axioms, it is useful in practice to divide the Tbox into a set of "definition" axioms and the remaining "general" axioms, which set is made as small as possible by applying a rewriting optimisation known as absorption. Tableau algorithms can then exploit the definitional part of the Tbox by using a lazy expansion technique.

. . . I think Jerome is right. Without the ONE-OF construct, DAML+OIL [what has now evolved to OWL] could be seen as a Tbox with "raw" RDF acting as an Abox.

I find two important ideas emerging at this point. First, the roles and purposes of the TBox and ABox are being made clearer, consistent with the definition we have been using [\[1\]](#) and with better clarity and applicability to the semantic Web than what was earlier presented in the *Description Logics Handbook* [\[3\]](#). And, second, the idea of a language split between OWL and RDF *viz* the TBox and ABox is made public for the first time.

Similarities to the Relational Model

The conceptual foundations of the [relational model](#) and RDF are indeed quite similar, based as they are on set theory and relationships. The relational model and description logics are both based on FOL ([first-order predicate logic](#)). The linkage with the data table of relational database systems is especially close and direct and was first a topic of a design architecture document by Tim Berners-Lee in 1998 [\[4\]](#).

Indeed, today:

- Many RDF triple stores are based on relational database (RDB) systems
- Close analogies exist between the RDB query language SQL and the RDF query language SPARQL
- There are effective working systems for translating from relational data to RDF (such as OpenLink's [RDF Views](#)), and
- A W3C incubator group is [now proposed](#) to transition into a full work group devoted to RDB to RDF mappings.

There is a rich literature to investigate these aspects in detail. And, of course, these matters are of critical import because 99% of the current structured data in the world is being managed by RDBMs [\[5\]](#).

However, of more direct interest to this specific series of articles is how this close relationship between RDB and RDF is viewed with respect to the ABox and TBox [separation of concerns](#).

Ian Horrocks (not surprisingly given the nature of his comments above), among others, has played a prominent role in looking at questions such as "hybrid DL-DB" (description logics + [relational] database) systems [\[6\]](#) and building conceptual links between relational databases and ontological-level reasoning [\[7\]](#).

We need not deconstruct his observations and arguments here in detail. What I glean from his strong background in description logics, however, is that relational data tables can be left *in situ* as ABox constructs, in the process gaining the efficiency of limited ABox reasoning and the efficiency of RDBMs. With proper design -- which I also understand to be pretty straightforward -- it should be possible to design hybrid ABox and TBox systems that work in a distributed context.

(Hmmm; sounds to me like ideas applicable to linked data !).

Horrocks' *et al.* more recent paper [\[7\]](#) and its expansion [\[8\]](#) propose an extension of OWL for such hybrid or split knowledge bases. These extensions are designed to allow modelers to designate a subset of TBox axioms as integrity constraints. For TBox-level reasoning, these axioms are treated as usual. However, these axioms may also be applied separately to ABox instance data to perform integrity checks. Integrity can then be checked in advance with those axioms ignored during standard TBox reasoning, thus also improving performance. I think these points also have direct relevance to linked data.

As a proponent of the OWL side of the spectrum, Horrocks' viewpoints have perhaps been too readily dismissed by some in the linked data community. Yet the major reason for looking at all of these questions from the perspective of description logics is to gain a [coherent](#) view across the entire semWeb enterprise. In the end, we are linking data for a *purpose*, to be able to do meaningful *work* with it. Just as

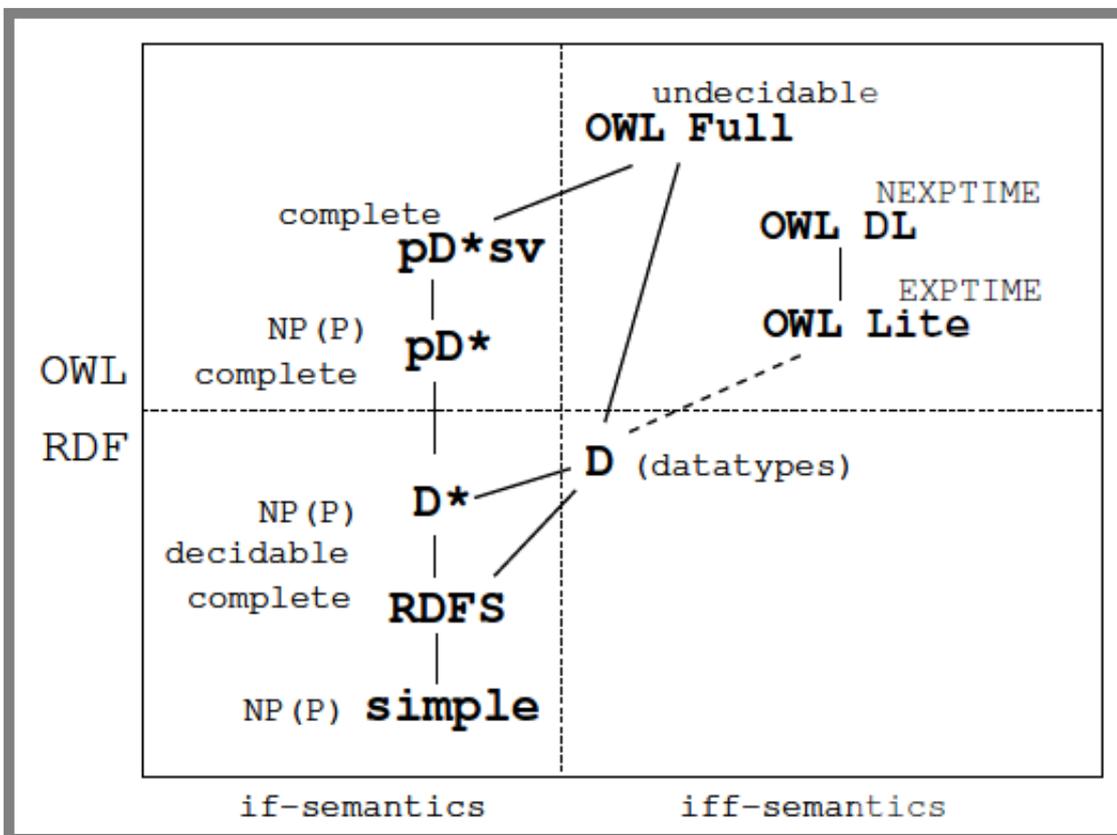
RDB data tables can be looked at and integrated productively as ABoxes in a DL construct, so may linked data.

What of a Simpler RDFS for the ABox?

If there truly is a separation of concerns between instance records (ABox) and reasoning constructs (TBox ontologies), what does that begin to tell us about the languages we need for these purposes? If we can postulate no OWL in a linked data instance construct (the ABox), why not narrow RDFS [9] as well in order to have a vocabulary only as expressive as what an instance record and its assertions require?

de Bruijn *et al* in 2005 demonstrated logically how RDF models can be related to description logics-based ontology languages, especially OWL DL, without the need to change syntax or semantics in either language [10]. They noted specifically the use of RDF graphs as ABoxes that could be readily queried using SPARQL.

Herman ter Horst wrote another influential paper in 2005 where he looked closely at the proofs of completeness and decidability for RDF and RDFS [11]. He defined a general RDF graph extension that was fully decidable, and importantly looked at each statement in the language from the standpoints of complexity and computational tractability. He was particularly seeking logics that would be more computationally efficient due to fewer entailments, while still being "decidable" (that is, provable to reach computational closure). Here is the basic chart of plotting the various language dialects he investigated:



He noted that inclusion of XML datatypes required the use of RDFS for closure and the addition of the so-called '*D*entailment*' could extend RDFS to include reasoning with datatypes. He then extended that

construct into what he called the '*pD*semantics*,' which was intended to allow useful conclusions to be drawn about instances in the presence of an ontology with relatively low computational complexity.

What this construct means -- as I understand it in the context of this series -- is that specialized dialects (*pD**) could govern the work of instance checking and other specialized work at the ABox level (RDFS) while being fully compatible with the TBox level (OWL) ontology. This means that languages and dialects could be tailored for the work at hand for efficiency and representational reasons, while maintaining logical integrity. Indeed, this very *pD** dialect of OWL is now included as one of the proposed profiles, OWL 2 RL, for the new release of OWL 2 [\[12\]](#).

In a different vein, the paper that won the best award at ESWC in 2007 looked at the question of simplifying RDFS [\[13\]](#). The authors were able to identify a fragment of RDFS that captured the complete semantics of RDF by carefully removing pieces that only described or allowed reasoning of the language itself.

A relatively streamlined and simplified structure for the ABox is not a new idea. Through version 3x, the [Protégé](#) ontology editor included a built-in for [SWRL](#) (the Semantic Web Rule Language) that included an ABox ontology [\[14\]](#):

```
<rdf:RDF xml:base='http://swrl.stanford.edu/ontologies/built-
ins/3.3/abox.owl'>
<owl:Ontology rdf:about=' '/>
<swrl:Builtin rdf:ID='hasValue' />
<swrl:Builtin rdf:ID='hasURI' />
<swrl:Builtin rdf:ID='isNumeric' />
<swrl:Builtin rdf:ID='notNumeric' />
<swrl:Builtin rdf:ID='isIndividual' />
<swrl:Builtin rdf:ID='isConstant' />
<swrl:Builtin rdf:ID='hasClass' />
<swrl:Builtin rdf:ID='hasProperty' />
<swrl:Builtin rdf:ID='hasIndividual'>
<swrlb:maxArgs rdf:datatype='http://www.w3.org/2001/XMLSchema#int'>1</
swrlb:maxArgs>
<swrlb:args rdf:datatype='http://www.w3.org/2001/XMLSchema#int'>1</swr
lb:args>
<swrlb:minArgs rdf:datatype='http://www.w3.org/2001/XMLSchema#int'>1</
swrlb:minArgs>
</swrl:Builtin>
<swrl:Builtin rdf:ID='setValue' />
</rdf:RDF>
```

Don't be fooled by the OWL designation in this file; for these uses, Protégé by convention requires all of its files to be of the OWL type. Note the simple vocabulary above has solely RDF predicates. We do not

think this is yet the correct design (see Part 4), but it captures the right idea.

Logical and mathematical advances since the first releases of RDF and OWL now suggest that, with proper care and design, various dialects or fragments can be designed for specific purposes and for computational efficiency while maintaining -- in their combination -- logical integrity. An RDFS fragment, if you will, dedicated for linked instance data and ABox instance record purposes, appears conceptually doable. And, it may be computationally advisable.

SAOR: One Approach to Combining These Pieces

The [SWSE](#) ("swizzy") project from [DERI](#) and the [National University of Ireland](#) in Galway has an interesting legacy and has been combining many of these threads into one approach to a semantic web search engine (hence, [SWSE](#)). You can use and test for yourself the new [VisiNav](#) interface and service, the newest instantiation of SWSE.

Relative to ABox (instance) data, the volume of TBox (structural) data on the Web is small: only around 0.7% of statements were classifiable as TBox statements [16].

In building what they call SAOR (for Scalable, Authoritative OWL Reasoner) for the SWSE effort, Aidan Hogan, Andreas Harth and Axel Polleres have intersected a number of interesting approaches and have taken some innovative paths to the questions of separating the TBox and ABox [15]. They have further applied this to the large-scale [Billion Triples Challenge](#) with interesting findings and results [16].

In their approach to building SAOR, the designers:

- Found, at scale, that complete inferencing at the instance level is neither feasible nor desirable
- Separated terminological knowledge (TBox) from assertional data (ABox) according to their use of the RDFS and OWL vocabulary
- Did not carry over owl : sameAs inferences to the TBox, with their rationale being that such an approach was in-line with a first-order-logic point of view where equalities do not affect predicates
- Stored all forward-chained inferences in the ABox only
- Tested for and removed what they called "ontology hijacking", which they defined as third parties attempting to broaden the definition of concepts for which they were not the original, authoritative authors, and
- Constructed a resulting synthesized, static TBox based on these rules.

They further picked up on a variant of the ter Horst pD*semantics noted above to speed the reasoner for calculating the forward-chaining inferences.

According to these decisions and rules, they found the overwhelming majority of statements within the 315,000 sources they crawled as being "non-authoritative" and indeed made many decisions that, in essence, threw out statements in the source instance sets. One interpretation, related to the thesis of this series but not directly noted by the authors, is that much of the linked data presently available on the Web is either over-specified or mis-specified. (I would argue that is due in part to linked data instance records trying to do more than their natural assertional role.)

Now, perhaps one could quibble with the rules and the decisions the authors employed (indeed, we do), but that is a topic for another day. What is interesting about the entire SAOR approach, I think, is its close attention to value and authoritativeness, all being split and recast into more tractable ABox and TBox portions, for handling reasoning at scale over large numbers of instances.

In my opinion, this is a seminal approach to the next generation of linked data that warrants much inspection and discussion.

Research on Other Work Tasks

A cursory discussion of the literature also shows some efforts that address the interstitial *work* areas noted in my conceptual architecture from [Part 2](#).

Full-text Search Engines

[Part 2](#) discussed full-text [search engines](#) in the broad semantic sense, and not specifically related to the ABox-TBox split. A couple of those and some others deserve a look because of their tighter integration of full-text search and attention to work splits.

A sister project to SWSE is [Sindice](#), which also uses Solr and employs the ter Horst pD* semantic framework [\[17\]](#). An inspection of [Sweet Tools](#), the semantic Web and -related tools listing, also suggests [Aperture](#) (a broadscale, full-text harvester with semantic capabilities); [LARQ](#) (which adds free text search to ARQ); [Virtuoso](#) (full-text and faceted search on top of a universal datastore); [Watson](#) (full-text search of metadata fields); and [Zebra](#) (specializing in structured library data and related).

Identity Relations

A couple of different approaches are being taken to identity testing, similarity or relations. The more direct approach is to do identity matching with a canonical ID or similar.

The SWSE group has one approach to object consolidation [\[18\]](#), which uses a clever method based on the `owl:InverseFunctionalProperty` (IFP) for performing large-scale consolidation of object identifiers for equivalent instances across data sources. Yet, as the authors note,

"One major issue we discovered involved `foaf:weblog`, which is defined as being an `owl:IFP` in the FOAF ontology. Thus its semantics mandates that the property uniquely defines the `foaf:Person` instance for which it is described."

This is both bad logic and wrong in many cases [see [19](#) for a critique]. The authors therefore needed to drop this assignment from their method. But, frankly, I think the broader problem again is too much predicate firepower for what should be a simple assertion that Joe Farmer has a blog (in fact, may have three!) and here are their URIs.

A very large, multi-year project to assign unique identifiers to entities is the [OKKAM](#) project [\[20\]](#). The intent is to provide a single and globally unique identifier for any entity through an '*Entity Name System*',

plus tools. Many methods will be employed to assign the identity relationship; specifics are still forthcoming with dozens of researchers working on the problem. I should note that the reference paper also touches upon some of the massive challenges associated with the current use of `owl:sameAs`.

Others have questioned a centralized ID service, instead preferring a mechanism that is more local and builds on co-reference research [21]. The [ReSIST](#) project has noted some of the issues of `owl:sameAs` use and management. It has proposed, instead, a '*Consistent Reference Service*' (or CRS). Asserting a co-reference in this approach is like its use in linguistics: it means a URI that describes the same entity, as does 'he', 'she' or 'it' as a co-reference in a sentence. This predicate indicates that the two resources are describing the same thing without carrying all of the heavy entailment of the `owl:sameAs` predicate that semantically means the two resources are exactly the same. The CRS are proposed to be set up and managed locally and in a distributed fashion.

A very different approach to identity assignments is Rough DL, which is a qualitative, "fuzzy" ontology for relating entities or concepts to their similar resources [22]. The method has also been applied to the very difficult problem of bibliographic records [23], where similarity is harder to judge because of use of initials and abbreviations. Rough DL may be especially appealing because even with the best state-of-the-art, there are error rates in any of the identity relating or disambiguation methods available. And, rather than try to assess these similarities with a probabilistic score, the "fuzzy" approach may even be one that can be reasoned over.

Disambiguation

To my knowledge, there is no disambiguation of entities presently taking place for distributed linked data sets. But, if not already, it soon will.

An example of how such a service might occur is the [uBio Taxonomic Name Server](#) from [The Marine Biological Laboratory at Woods Hole](#). Via Web service or direct HTML form, an entity name (in this case a biological species) or its variants can be submitted for disambiguation and assignment to the proper identifier (name).

There is much research behind the algorithms and approaches to named entity disambiguation beyond the scope of this present series.

Some Concluding 'Big Picture' Implications

Our arguments to this point in this series do not suggest nor require that current practice need change. Clearly, we are seeing growth, uptake and use with current practices regarding linked data.

One of the beautiful aspects of RDF as a data model and the semantic Web is that the underlying languages and standards are so flexible. Find a way to do stuff better in the future? Fine; go ahead and do it, because what has come before can be easily transitioned or accommodated.

The real thrust of this series has been "best practice." There are certainly many viewpoints on that topic, and the understanding of it for a linked data environment at scale is also evolving. This is healthy, vibrant and exciting. Who knows what is truly best practice? I personally believe the market will determine that

by what gets adopted and becomes self-sustaining by providing value.

However, as [Structured Dynamics](#) attempts to think through these issues -- to look seriously at moving from simply proving the exposure of Web data to one of meaningfully doing *work* and relying on it at large scale -- we see warts and challenges. Such is growth. It is natural. And change does not mean that what came before was wrong.

So, what do we see as some of these 'big picture' implications?:

- Linked data is surely proving the idea of a Web of Data and is bringing broader awareness about the usefulness of the semantic Web, its languages, and its standards
- The semantic Web does not require "reasoning" at the point of linked data publishing, only that the linked data be published in a form that *ultimately* supports reasoning
- Linked data is largely a basis for exposing and sharing instance data
- As instance records, linked data is about assertions and attributes, not *provability* or *decidability*
- Linked data can be guided by the constructs of description logics
- Linked data instance records only require a subset of RDFS to be structured sufficiently for assertions and attributes
- Work applied to linked data in a distributed dataset setting can be segregated and optimized
- This specialized work (*identity testing*, *disambiguation* or *full-text retrieval*) does not belong at the ABox level, but also can be conducted separately from standard TBox reasoning and inference
- Linked data has appropriately embraced RDF, but has often overstepped its natural bounds by "cherry-picking" OWL predicates without regard to actual use in an open-world knowledge base
- As linked data is incorporated into knowledge bases, fragments and dialects of both RDF and OWL can be applied to specific work tasks to improve scalability and computational tractability
- `owl:sameAs` and other `owl:statements` within linked data instance records are being rejected by aggregation and consolidation services; let's figure better ways to assert identities, memberships and linkages without entailing what is not being (and can not logically be) supported.

If we can do these things, we can simplify what it means to publish "linked data-ready" structured data. Being *coherent* about these matters is a key.

[1] This is our [working definition](#) for [description logics](#):

"Description logics and their semantics traditionally split *concepts* and their relationships from the different treatment of *instances* and their attributes and roles, expressed as fact assertions. The concept split is known as the TBox (for *terminological* knowledge, the basis for *T* in *TBox*) and represents the schema or taxonomy of the domain at hand. The TBox is the structural and intensional component of conceptual relationships. The second split of instances is known as the ABox (for *assertions*, the basis for *A* in *ABox*) and describes the attributes of instances (and individuals), the roles between instances, and other assertions about instances regarding their class membership with the TBox concepts."

[2] Patrick Hayes, 2004. "RDF Semantics," a *W3C Recommendation*, February 2004. See [http://www.w3.org/TR/rdf-
mt/](http://www.w3.org/TR/rdf-mt/).

[3] F. Baader, D. Calvanese, D. McGuinness, D. Nardi, and P. F. Patel-Schneider, editors. *The Description Logic Handbook: Theory, Implementation and Applications*. Cambridge University Press, 2003. Sample chapters may be viewed from Enrico Franconi's Description Logics course notes and tutorial at <http://www.inf.unibz.it/~franconi/dl/course/>, which is an excellent starting reference point on the subject.

[4] See Tim Berners-Lee, 1998. *Relational Databases on the Semantic Web*; see <http://www.w3.org/DesignIssues/RDB-RDF.html>; updated in 2002. (An interesting aside from that document was its observation that there may only be a limited number of RDF ontologies moving forward, with quite the opposite now apparent.) For a general discussion of relationships between RDF and standard relational data tables, I recommend: Andrew Newman, 2007. "A Relational View of the Semantic Web," an online article at *XML.com*, March 14, 2007; see <http://www.xml.com/pub/a/2007/03/14/a-relational-view-of-the-semantic-web.html>.

[5] Yet structured data is typically seen as only a minor portion of all available information, with text (unstructured) data generally estimated as comprising 80% to 85% of all useful information. Most text information is managed by text engines [also known as search or information retrieval (IR) index engines], and not by RDBMs.

[6] Sean Bechhofer, Ian Horrocks, and Daniele Turi. "The OWL Instance Store: System Description," in *Proceedings of the 20th International Conference on Automated Deduction (CADE-20)*, Lecture Notes in Artificial Intelligence, pages 177-181. Springer, 2005. See <http://www.comlab.ox.ac.uk/people/ian.horrocks/Publications/download/2005/BeHT05.pdf>.

[7] Boris Motik, Ian Horrocks, and Ulrike Sattler, 2007. "Bridging the Gap Between OWL and Relational Databases," in *Proceedings of the Sixteenth International World Wide Web Conference (WWW 2007)*. See <http://www.comlab.ox.ac.uk/people/ian.horrocks/Publications/download/2007/MoHS07a.pdf>.

[8] Boris Motik, Ian Horrocks and Ulrike Sattler, 2006. *Integrating Description Logics and Relational Databases*, Technical Report from the University of Manchester, UK, December 6, 2006, 44 pp. See <http://web.comlab.ox.ac.uk/people/Boris.Motik/pubs/mhs06constraints-report.pdf>. This is really a very in-depth, comprehensive treatment of the subject.

[9] We are still evaluating whether the simpler ABox vocabulary discussed in Part 4 can be met solely with RDF or requires RDFS. Our preference is RDF, but subClassOf relationships and reasoning over datatypes [11] seem to point to RDFS. What is presented in Part 4 is still quite preliminary and subject to change.

[10] Jos de Bruijn, Enrico Franconi and Sergio Tessaris, 2005. "Logical Reconstruction of Normative RDF," presented at the *International Workshop on OWL: Experiences and Directions (OWLED 2005)*, Galway, Ireland. See <http://www.inf.unibz.it/~jdebruijn/publications/owl-05.pdf>.

[11] Herman J. ter Horst, 2005. "Completeness, Decidability and Complexity of Entailment for RDF Schema and a Semantic Extension involving the OWL Vocabulary," *Journal of Web Semantics*, Vol. 3, 2005, pp. 79-115. See <http://www.websemanticsjournal.org/papers/20050719/document5.pdf>.

[12] Boris Motik et al., eds., 2008. "OWL 2 Web Ontology Language: Profiles," a *W3C Working Draft*, December 2, 2008. See <http://www.w3.org/TR/owl2-profiles/>, esp. http://www.w3.org/TR/owl2-profiles/#OWL_2_RL. There are also reasoners available for dealing with such specialized dialects, such as for example, <http://www.ivan-herman.net/Misc/PythonStuff/RDFSClosure/Doc/RDFSClosure-module.html>.

[13] Sergio Muñoz, Jorge Perez and Claudio Gutierrez, 2007. "Minimal Deductive Systems for RDF," paper presented at the *4th European Semantic Web Conference, ESWC 2007*, Innsbruck, Austria, June 2007. See <http://www2.ing.puc.cl/~jperez/papers/minimal-rdf-camera-ready-ext.pdf>.

- [14] See <http://protege.cim3.net/cgi-bin/wiki.pl?SWRLABoxBuiltIns> for a description; the actual ontology may be found at <http://swrl.stanford.edu/ontologies/built-ins/3.3/abox.owl>.
- [15] Aidan Hogan, Andreas Harth and Axel Polleres, 2008. "SAOR: Authoritative Reasoning for the Web," in *Proceedings of the 3rd Asian Semantic Web Conference (ASWC 2008)*, Bangkok, Thailand, Dec. 2008. See <http://sw.deri.org/~aidanh/docs/aswc08.pdf>.
- [16] Aidan Hogan, Andreas Harth and Axel Polleres, 2008. "Scalable Authoritative OWL Reasoning on a Billion Triples," in *Proceedings of Billion Triple Semantic Web Challenge 2008*, at the *7th International Semantic Web Conference (ISWC2008)*, Karlsruhe, Germany, 2008. See http://sw.deri.org/~aidanh/docs/saor_billiontc08.pdf.
- [17] Giovanni Tummarello, Renaud Delbru and Eyal Oren, 2007. "Sindice.com: Weaving the Open Linked Data," in *Proceedings of the International Semantic Web Conference (ISWC), 2007*. See presentation slides at <http://www.eyaloren.org/slides/iswc2007.pdf>.
- [18] Aidan Hogan, Andreas Harth and Stefan Decker, 2007. "Performing Object Consolidation on the Semantic Web Data Graph," in *1st I3 Workshop: Identity, Identifiers, Identification Workshop*, 2007. <http://sw.deri.org/2007/02/objcon/paper.pdf>.
- [19] See <http://esw.w3.org/topic/InverseFunctionalProperty>.
- [20] Paolo Bouquet, Heiko Stoermer, Claudia Niederee, and Antonio Mana, 2008. "Entity Name System: The Backbone of an Open and Scalable Web of Data," in *Proceedings of the IEEE International Conference on Semantic Computing, ICSC 2008*, August 2008. See <http://www.okkam.org/publications/stoermer-EntityNameSystem.pdf>.
- [21] Afraz Jaffri, Hugh Glaser and Ian Millard, 2008. "URI Identity Management for Semantic Web Data Integration and Linkage," in *3rd International Workshop On Scalable Semantic Web Knowledge Base Systems*, November 25-30, 2007, Vilamoura, Algarve, Portugal. See http://eprints.ecs.soton.ac.uk/14361/1/URI_Identity_Management_for_Semantic_Web_Data_Integration_and_Linkage.pdf.
- [22] Stefan Schlobach, Michel Klein and Linda Peelen, 2007. "Description Logics with Approximate Definitions: Precise Modeling of Vague Concepts," in *Proceedings of the 20th International Joint Conference on Artificial Intelligence, IJCAI 07*, Hyderabad, India, January 2007; see <http://www.ijcai.org/papers07/Papers/IJCAI07-088.pdf>.
- [23] Michel C.A. Klein, Peter Mika and Stefan Schlobach, 2007. "Rough Description Logics for Modeling Uncertainty in Instance Unification," in *Proceedings of the Third ISWC Workshop on Uncertainty Reasoning for the Semantic Web*, Busan, Korea, November 12, 2007. See <http://ftp.informatik.rwth-aachen.de/Publications/CEUR-WS/Vol-327/paper4.pdf>.