



Stream Reasoning – Research and Use cases at the Institute of Information Business

Axel Polleres & Javier Fernandez, WU Wien

Outline

Disclaimer: **NO concrete work on Stream Reasoning in my group at the moment**, but still on 3 related topics...

1. Dynamic Linked Data & Open Data
2. Efficient RDF Stream Interchange (ERI)
3. SPARQL 1.1 Update & Entailment as a basis for Stream Reasoning Semantics?

1. Dynamic Linked Data & Open Data:

Collecting changes in Linked and Open Data:

- Dynamic Linked Data Observatory
- Open Data Portalwatch
- The DBpedia Wayback machine

1.1 Dynamic Linked Data Observatory

Started in DERI since 2012, **weekly snapshots** of crawls of Linked Data... Interesting dataset for investigating LOD dynamics.

<http://dyldo.deri.org/>

Publications:

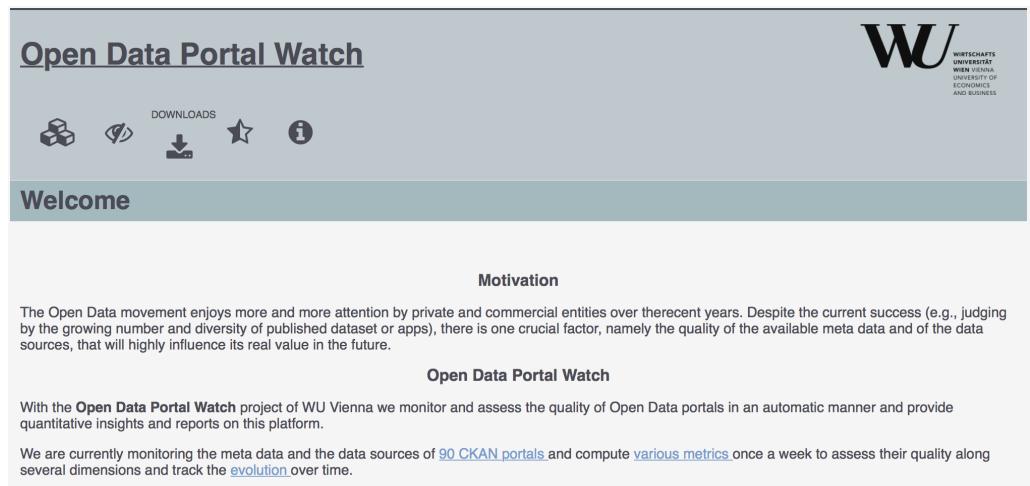
Tobias Käfer, Ahmed Abdelrahman, Jürgen Umbrich, Patrick O'Byrne, Aidan Hogan. "**Observing Linked Data Dynamics**". In the Proceedings of the 10th Extended Semantic Web Conference (ESWC 2013), Montpellier, France, 26–30 May, 2013.

Tobias Käfer, Jürgen Umbrich, Aidan Hogan and Axel Polleres, "**Towards a Dynamic Linked Data Observatory**", In the Proceedings of the Linked Data on the Web WWW2012 Workshop ([LDOW 2012](#)), Lyon, France, 16 April, 2012.

1.2 OPEN DATA PORTAL WATCH

<http://data.wu.ac.at/portalwatch/>

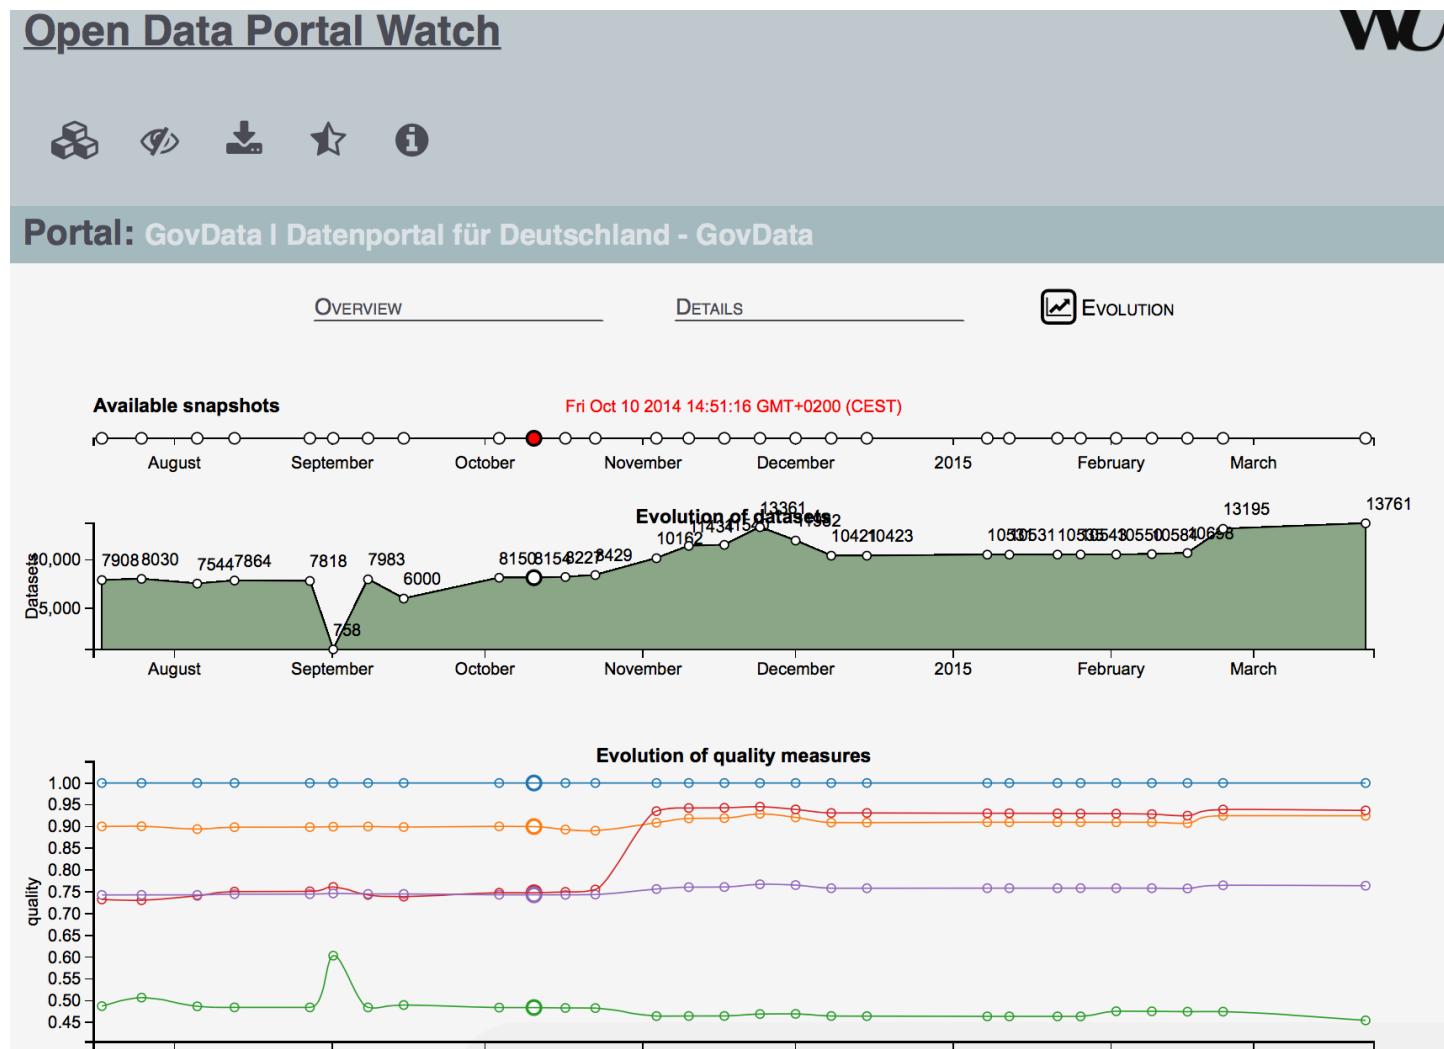
- Started in 2014 – similar idea, different dataset:
 - Weekly crawls of Data and Metadata from over 90 Open Data Portals (CKAN, Socrata, etc)
- Goal: Quality assessment
- Evolution tracking
 - Meta data
 - Data



The screenshot shows the homepage of the Open Data Portal Watch. At the top, there's a navigation bar with icons for data cubes, eye, download, star, and info. Below it is a green header bar with the text "Welcome". The main content area has a title "Open Data Portal Watch" and a "Motivation" section. The motivation text explains that the Open Data movement is growing and discusses the importance of quality meta data. It also mentions the project's goal of monitoring and assessing portal quality. At the bottom, there's a note about monitoring data sources and computing metrics.

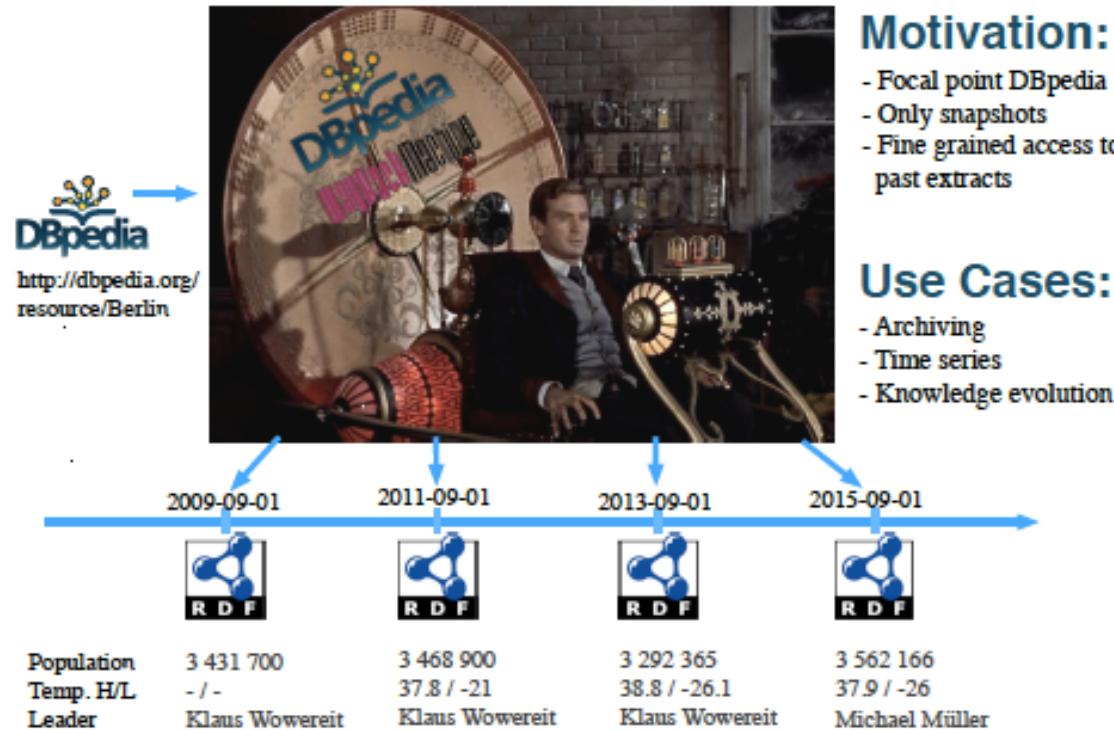
- Jürgen Umbrich, Sebastian Neumaier, and Axel Polleres.
 Quality assessment & evolution of open data portals. In *IEEE International Conference on Open and Big Data*, Rome, Italy, August 2015. **Best paper award**.

1.2 Open Data Portal Watch



Dbpedia Wayback Machine

Travelling DBpedia back in time!



Motivation:

- Focal point DBpedia
- Only snapshots
- Fine grained access to past extracts

Use Cases:

- Archiving
- Time series
- Knowledge evolution

<http://data.wu.ac.at/wayback/> ... Different concept: create historical RDF Data from wikipedia/dbpedia revision history

"The DBpedia wayback machine" by Javier Fernandez, Patrik Schneider, Jürgen Umbrich, September 2015, SEMANTiCS2015 (Best Poster).

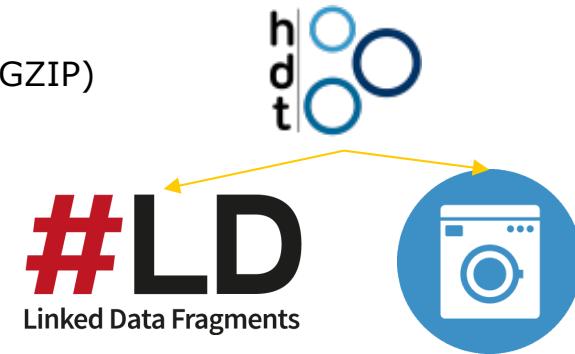
1. Dynamic Linked Data & Open Data – Discussion:

Stream Reasoning related questions:

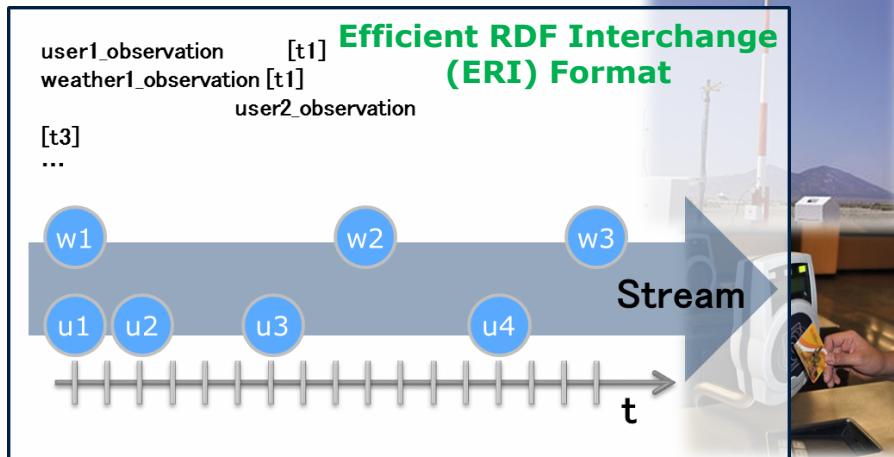
- How can we efficiently store and query archival data?
- Where does Reasoning help us? (e.g. detect anomalies, conflicts, assume defaults in case of missing crawls, etc.)
- Javier D. Fernández, Axel Polleres, and Jürgen Umbrich. Towards efficient archiving of dynamic linked open data. In *Managing the Evolution and Preservation of the Data Web - First Diachron Workshop at ESWC 2015*, pages 34-49, Portorož, Slovenia, May 2015

2. Efficient RDF Stream Interchange (ERI)

- Lightweight **Binary RDF (HDT)**
 - Highly **compact** serialization of RDF (slightly bigger than GZIP)
 - compact RDF store (without prior decompression)
 - Basics for successfully projects
 - Linked Data fragments
 - LOD Laundromat



- **RDF Data Streams**



Querying & Efficient Serialization of streams of LD

- In the light of IoT/Industry 4.0, etc. ...



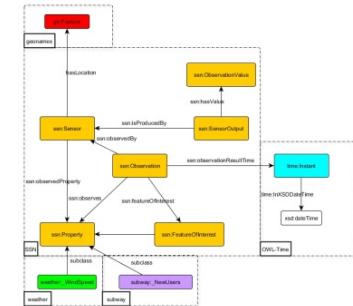
- We are monitoring the W3C RSP (RDF Stream Processing) CG, but also EXI (Efficient XML Interchange) and WoT WG

Efficient Serialization of streams of LD

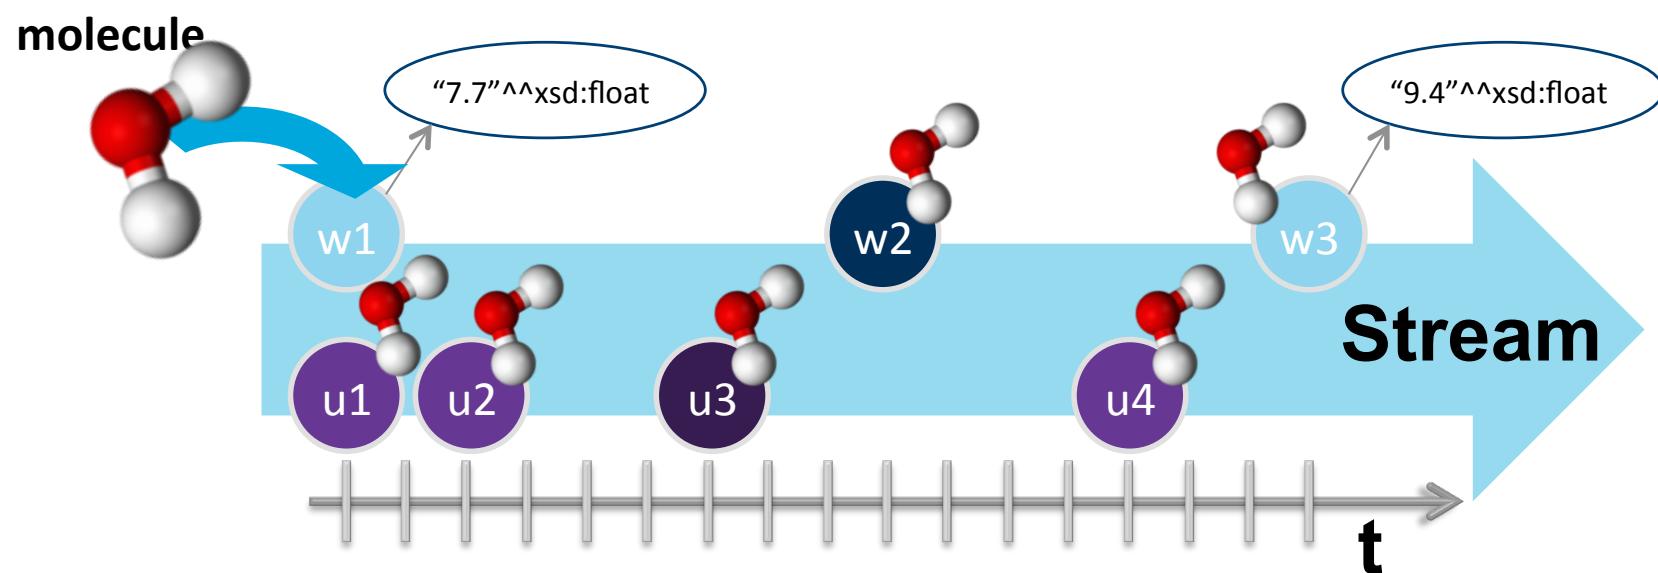
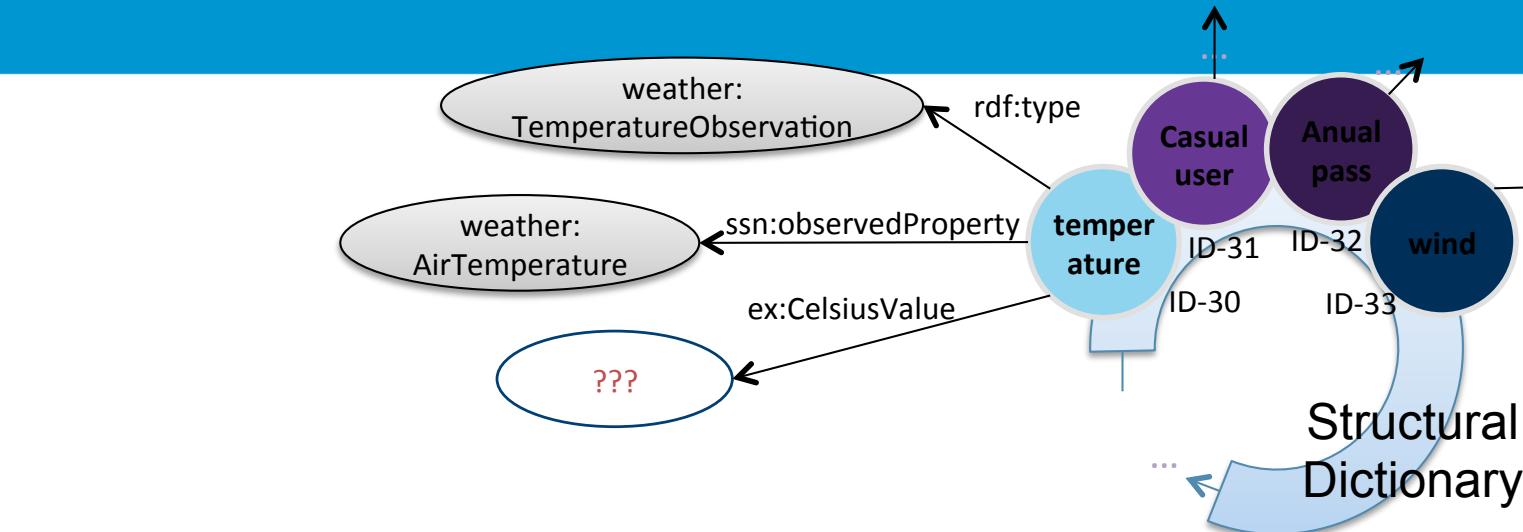
	Plain: Turtle/ Trig/ JSON-LD	Plain +Compressio n (e.g. gzip)	HDT	Streamin g HDT	RDSZ	RDF/ XML + EXI	ER I
Streamable	Yes	Yes	No	Yes	Yes	Yes	Yes
Scalable	Limited	Yes	Yes	No	Yes	Yes	Yes
Easy (fast) to create and parse	Yes	Limited	Limited	Yes	Limited	Limited	Yes
Compact	No	Yes	Yes	Limited	Yes	Yes	Yes
Parametrizable : compression/time	No	Limited	Yes	No	Limited	Limited	Yes

EFFICIENT RDF INTERCHANGE (ERI) FORMAT – Basic Concepts

- (Assumption) Most RDF streams are **well structured**
 - the **structure** is well-known by the data provider
 - the number of **variations** in the structure are limited
- Efficient RDF Interchange (ERI) Format encodes the information at two levels:
 - A sliding dictionary of **structures**: Structural Dictionary
 - The **concrete value** for each predicate



EFFICIENT RDF INTERCHANGE (ERI) FORMAT – Basic Concepts



2. Efficient RDF Stream Interchange (ERI)– Discussion:

Stream Reasoning related questions:

- Where does Reasoning help us? (e.g. compact representation by stripping off implicit information)
- Vague idea what could be related here:

*Reinhard Pichler, Axel Polleres, Sebastian Skritek, and Stefan Woltran.
Complexity of redundancy detection on rdf graphs in the presence of rules,
constraints, and queries. Semantic Web - Interoperability, Usability,
Applicability (SWJ), 4(4), 2013.*

3. SPARQL 1.1 Update & Entailment as a basis for Stream Reasoning Semantics?

SPARQL 1.1 can update a stream and do entailment regimes, however...

Standardization of SPARQL 1.1 Update, and SPARQL 1.1 Entailment Regimes with triple stores implementing those standards (rewriting- or materialization-based) **is silent about their interaction**

What's been done already?

- What do off-the-shelf triple stores do?
 - **Entailment** typically handled
 - at (bulk) loading by ***materialization*** , or
 - at *query time* by ***rewriting***
but **not in the context of Updates.**
 - no “standard” behavior for **Delete & Insert** upon materialized stores.
 - interplay of Entailments and Update left out in the SPARQL 1.1 spec.
- What does the literature say?

Approaches *in the literature on updates and RDFS* (or also DLs) limited to ***atomic update*** operations...

- [Gutierrez et al., ESWC2011] ABox deletions in the context of RDFS
- [Calvanese et al., ISWC2010] ABox & TBox insertions in the context of DL-Lite (incl. inconsistency repair)

Also related:

- Deductive DBs: [Gupta et al., SIGMOD93]: **DRed (delete and re-derive)**, applied by [Kotowski et al.2011] and [Urbani et al. 2013] in the context of RDF/RDFS...
- KB evolution, Belief revision, etc.: Various works in classical AI and philosophy

Particularly, none of these considers the interplay between DELETE, INSERT based on a joint WHERE clause as in SPARQL

Our initial thoughts on this problem...



Updating RDFS ABoxes and TBoxes in SPARQL

Albin Ahmeti¹, Diego Calvanese², and Axel Polleres³

¹ Vienna University of Technology, Favoritenstraße 9, 1040 Vienna, Austria

² Faculty of Computer Science, Free University of Bozen-Bolzano, Bolzano, Italy

³ Vienna University of Economics and Business, Welthandelsplatz 1, 1020 Vienna, Austria

Abstract. Updates in RDF stores have recently been standardised in the SPARQL 1.1 Update specification. However, computing entailed answers by ontologies is usually treated orthogonally to updates in triple stores. Even the W3C SPARQL 1.1 Update and SPARQL 1.1 Entailment Regimes specifications explicitly exclude a standard behaviour for entailment regimes other than simple entailment in the context of updates. In this paper, we take a first step to close this gap. We define a fragment of SPARQL basic graph patterns corresponding to (the RDFS fragment of) *DL-Lite* and the corresponding SPARQL update language, dealing with updates both of ABox and of TBox statements. We discuss possible semantics along with potential strategies for implementing them. In particular, we treat both, (i) materialised RDF stores, which store all entailed triples explicitly, and (ii) reduced RDF Stores, that is, redundancy-free RDF stores that do not store any RDF triples (corresponding to *DL-Lite* ABox statements) entailed by others already. We have implemented all semantics prototypically on top of an off-the-shelf triple store and present some indications on practical feasibility.

3. SPARQL 1.1 Update & Entailment as a basis for Stream Reasoning Semantics? – Discussion:

- A first step to close the gap left by the current standards (**SPARQL1.1 Update** vs. **SPARQL1.1 Entailment Regimes**)
 - Seemingly no “one-size fits all” semantics: Particularly, DRed not intuitive in all cases!
 - Query rewriting may in some cases be more efficient than re-materialization
 - Non-intuitive corner cases in **each possible** semantics
→ depends on use case?

Which stream scenarios scenarios warrant which semantics?