



WIRTSCHAFTS
UNIVERSITÄT
WIEN VIENNA
UNIVERSITY OF
ECONOMICS
AND BUSINESS



Handling Inconsistencies due to Class Disjointness in SPARQL Updates

joint work with: Albin Ahmeti, Diego Calvanese, Vadim Savenkov

Axel Polleres

web: <http://polleres.net>

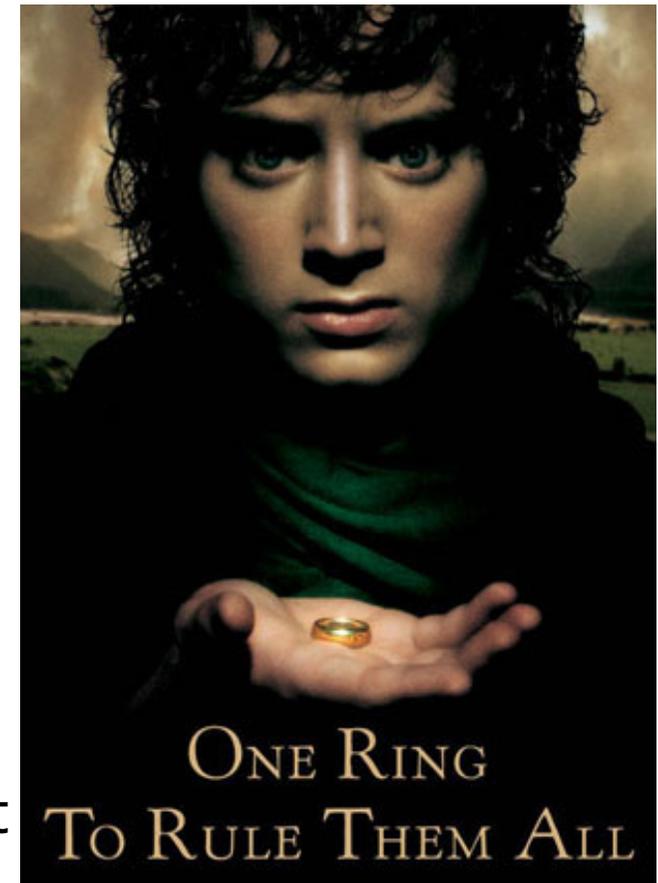
twitter: @AxelPolleres

The quest...

One Semantics to rule them all

SPARQL1.1 Updates and Entailment -
*Why the specification is silent about
their interaction*

- **SPARQL1.1 Update** allows to update RDF Graphs
- **SPARQL1.1 Entailment Regimes** tells us what answers a SPARQL query gives us including implicit triples
- But: What does it mean to update implicit triples?
- *Particularly (in this paper): How to deal with inconsistencies?*



In which **library** do I find the “Lord of the Rings” and where do I find it?

Wikipedia article for the National Library of Colombia. The article title is "National Library of Colombia". It includes a table of contents with sections: 1 Founding and history, 2 Mission, 3 Functions, 4 Law of legal deposit, 5 Catalog, 6 External links. The "Founding and history" section states: "The National Library of Colombia is generally considered to be the oldest national library in the Americas. It was founded at the end of the eighteenth century in 1777 by Viceroy".

DBpedia page for the National Library of Colombia. The page title is "About: National Library of Colombia". It displays a table of properties and their values:

S	P	O
:NLC	geo:long	4.609553
:NLC	geo:lat	-74.068649
:NLC	a	:Library
:NLC	:locatedIn	:Colombia
:LordOfTheRings	:inCatalogOf	:NLC

Additional properties shown below the table:

- dbpedia:WikiPageRevisionID: 654723423 (xsd:integer)
- dbpedia:director: Consuelo Gaitán Gaitán (en)
- dbpedia:established: 1777-01-09 (xsd:date)
- dbpedia:libraryLogo: dbpedia:File:LogoBNC.jpg
- dbpedia:libraryName: Biblioteca Nacional de Colombia (en)

SPARQL 1.1 Query language

- SPARQL offers a standard protocol/service interface to data offering services like DBPedia!

```
SELECT ?L WHERE {  
  :LordOfTheRings :inCatalogOf ?L.  
  ?L a :Library .  
}
```



- SPARQL

→ QU

← QU

L
http://dbpedia.org/resource/National Library of Colombia
http://dbpedia.org/resource/Le%C3%B3n de Greiff Library
http://dbpedia.org/resource/Spain Library (Medellín)
http://dbpedia.org/resource/León de Greiff Library

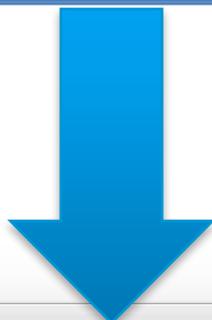
ON, ...)

SPARQL1.1 Entailment Regimes:

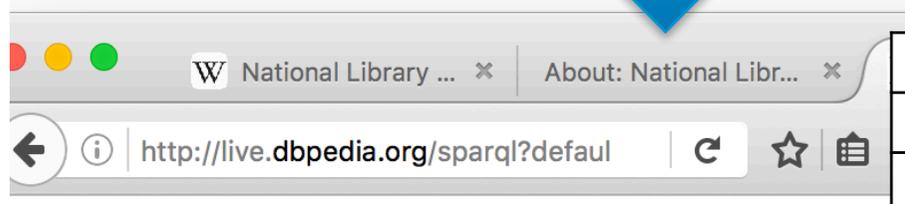
- Make use of ontological inferences (**RDFS** and **OWL**):

```

SELECT ?L WHERE {
  :LordOfTheRings :inCatalogOf ?L.
  ?L a schema:Library .
}
    
```



S	P	O
:NLC	geo:long	4.609553
:NLC	geo:lat	-74.068649
:NLC	a	:Library
:NLC	:locatedIn	:Colombis
:LordOfTheRings	:inCatalogOf	:NLC

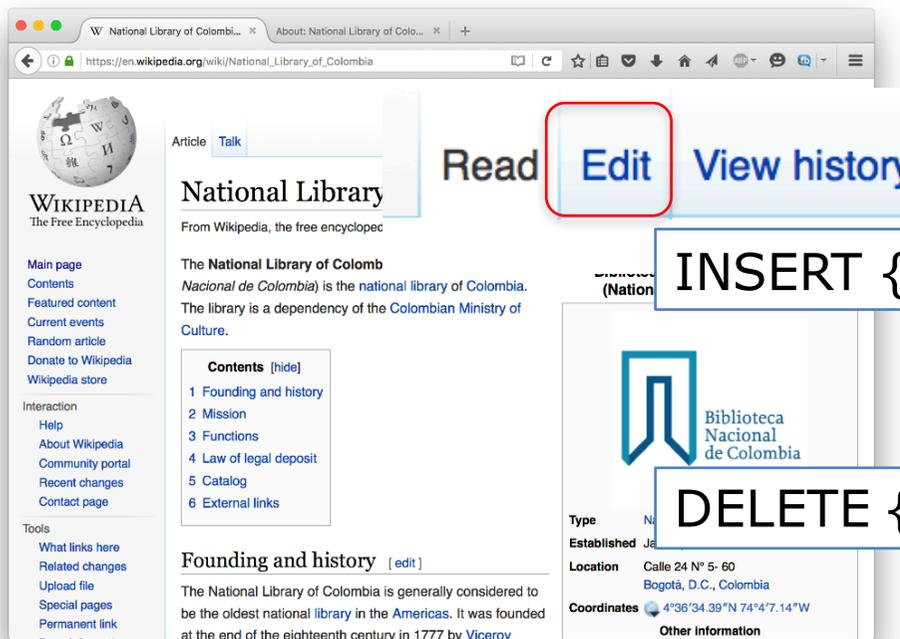


S	P	O
:Library	rdfs:subClassOf	:Organisation
:Library	rdfs:subClassOf	:schema:Library

L
http://dbpedia.org/resource/National Library of Colombia
http://dbpedia.org/resource/Le%C3%B3n de Greiff Library
http://dbpedia.org/resource/Spain Library (Medellín)

But: the Semantic Web is all about updates! SPARQL 1.1 Update

- What do updates mean?



```
INSERT { :NLC :a :Place. }
```

```
DELETE { :NLC a Or . . . }
```

```
DELETE { :NLC a :Organisation. }  
INSERT { :NLC :a Place. }  
WHERE { ?X geo:lat [ ] ; geo:long [ ] . }
```

What do these updates mean in an RDF store that does Entailment?

Previous work: What happened before...

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

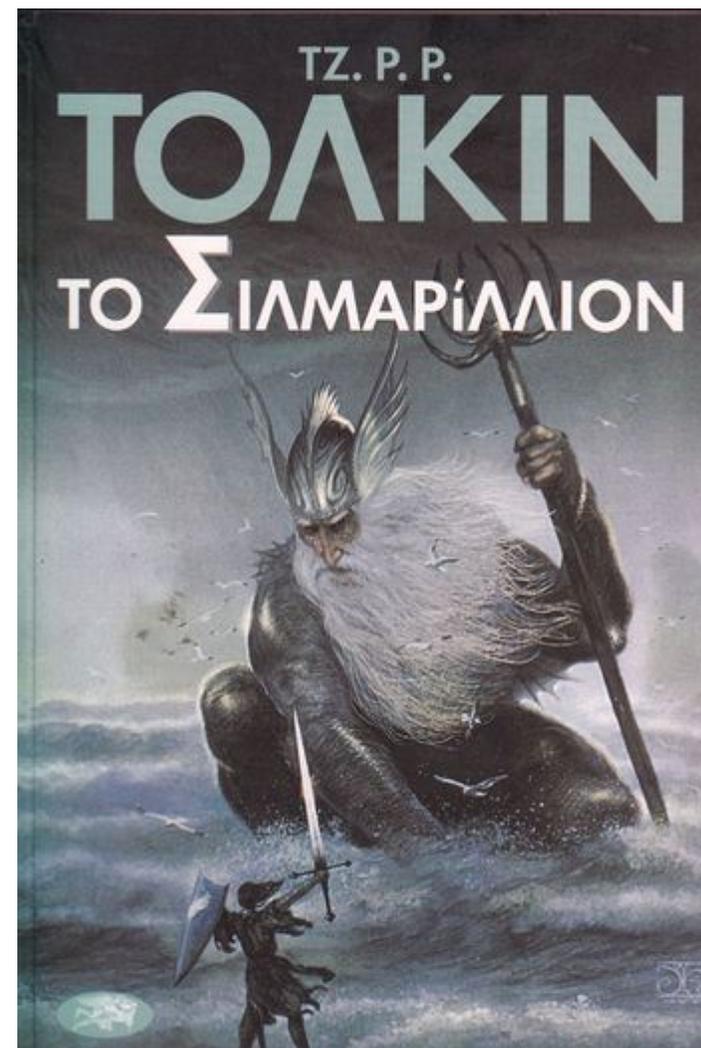
² Faculty of Computer Science, Free University of Bozen-Bolzano, Bolzano, Italy

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

Discussed several possible semantics for SPARQL Update under RDFS Entailment

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



Previous work: Our initial assumptions...

- Materialised store...

S	P	O	S	P	O
:NLC	geo:long	4.609553	:Library	rdfs:subClassOf	:Organisation
:NLC	geo:lat	-74.068649	:Library	rdfs:subClassOf	:schema:Library
:NLC	a	:Library			
:NLC	:hasDirector	:ConsueloGaitánGaitán			
:LordOfTheRings	:inCatalogOf	:NLC			
:NLC	a	:Organisation			
:NLC	a	schema:Library			

- Low expressivity ontology language... RDFS
- Semantics for update should:
 - Preserve materialisation
 - Not "leave traces"

Previous work: Our initial solution...

- Idea: keep Materialised state by rewriting updates:

DELETE { P_d }
INSERT { P_i }
WHERE { P_w }

- **Sem_2^{mat}**

- **Insert** the instantiations of P_i **plus all their effects.**
- **Delete** the instantiations of P_d **plus all their causes;**

$$G_u^{Sem_2^{mat}}(P_d, P_i, P_w) = G_u(P_d^{caus}, P_i^{eff}, \{P_w\} \{P_d^{fvars}\})$$

$$P_d^{fvars} = \{?x \text{ a rdfs:Resource.} \mid \text{for each } ?x \in Var(P_d^{causP_d}) \setminus Var(P_d)\}$$

Our initial solution... Example:

```
DELETE { ?X :a :Organisation. }
INSERT { ?X :a Place. }
WHERE { ?X geo:lat [] ; geo:long []. }
```

rewrite(u,T)

```
DELETE { ?X :a :Organisation. ?X a Library }
INSERT { ?X :a Place. ?X a schema:Place }
WHERE { ?X geo:lat [] ; geo:long []. }
```

TBOX

S	P	O
:Library	rdfs:subClassOf	:Organisation
:Library	rdfs:subClassOf	schema:Library
:Place	rdfs:SubClassOf	schema:Place

ABOX

S	P	O
:NLC	geo:long	4.609553
:NLC	geo:lat	-74.068649
:NLC	a	:Library
:NLC	:hasDirector	:ConsueloGaitánGaitán
:LordOfTheRings	:inCatalogOf	:NLC
:NLC	a	:Organisation
:NLC	a	schema:Library

Our initial solution... Example:

```
DELETE { ?X :a :Organisation. }
INSERT { ?X :a Place. }
WHERE { ?X geo:lat [] ; geo:long []. }
```

rewrite(u,T)

```
DELETE { ?X :a :Organisation. ?X a Library }
INSERT { ?X :a Place. ?X a schema:Place }
WHERE { ?X geo:lat [] ; geo:long []. }
```

TBOX

S	P	O
:Library	rdfs:subClassOf	:Organisation
:Library	rdfs:subClassOf	schema:Library
:Place	rdfs:SubClassOf	schema:Place

ABOX

S	P	O
:NLC	geo:long	4.609553
:NLC	geo:lat	-74.068649
:NLC	a	:Library
:NLC	:hasDirector	:ConsueloGaitánGaitán
:LordOfTheRings	:inCatalogOf	:NLC
:NLC	a	:Organisation
:NLC	a	schema:Library
:NLC	a	:Place
:NLC	a	schema:Place

for many use cases the most 'reasonable', among the semantics we looked into...

Let's revisit our initial assumptions...

- Materialised store...
 - ... fits e.g. DBpedia (all Abox inferences are materialised)
 - **consistent**
- Low expressivity ontology language... RDFS
 - ... does not quite fit DBpedia:
 - "OWL Dbpedia" :
 - `rdfs:subClassOf`, `rdfs:subPropertyOf` `rdfs:domain`, `rdfs:range`,
`owl:inverseOf`, **`owl:disjointWith`**
- Semantics for update should:
 - Preserve materialisation
 - Not "leave traces"
 - **Preserve consistency**

Inconsistencies!

Inconsistencies in DBPedia:

unfortunately **there are** inconsistencies in DBpedia...



S	P	O
:NLC	geo:long	4.609553
:NLC	geo:lat	-74.068649
:NLC	a	:Library
:NLC	:hasDirector	:ConsueloGaitánGaitán
:LordOfTheRings	:inCatalogOf	:NLC
:NLC	a	:Organisation
:NLC	a	schema:Library
:NLC	a	:Place

S	P	O
:Library	rdfs:subClassOf	:Organisation
:Library	rdfs:subClassOf	schema:Library
:Place	rdfs:SubClassOf	schemaPlace
:Place	owl:disjointWith	:Organisation

Can be introduced due to uncautious **updates** and the flexibility of mappings ☹

So, how can we do SPARQL updates that preserve consistency (and materialization?)

- Dealing with different forms of inconsistencies:
 - Intrinsic inconsistencies “within” updates
 - Inconsistencies between “old” and “new” knowledge
 - ... Different solution strategies:
 - Brave
 - Cautious
 - Fainthearted (somewhere in between ;-))

So, how can we do SPARQL updates that preserve consistency (and materialization?)

- Dealing with different forms of inconsistencies:
 - **Intrinsic inconsistencies “within” updates**
 - ... **solution: “safe” rewriting**
 - Inconsistencies between “old” and “new” knowledge
 - ... Different solution strategies:
 - Brave
 - Cautious
 - Fainthearted (somewhere in between ;-)

SPARQL updates: deal with inconsistency within *new knowledge*

`:Place owl:disjointWith :Organisation .` $Place \sqsubseteq \neg Organisation$
`:based_near rdfs:domain :Organisation.` $\exists based_near \sqsubseteq Organisation$
`:based_near rdfs:range :Place.` $\exists based_near^- \sqsubseteq Place$

"Unsafe" update → intrinsically inconsistent:

```
INSERT {?X :based_near ?Y } WHERE { ?X :locatedIn ?Y .}
```

`:NLC :locatedIn :Bogotá .`

`:Bogotá :locatedIn Colombia .`

SPARQL updates: deal with inconsistency within *new knowledge*

:Place owl:disjointWith :Organisation .

:based_near rdfs:domain :Organisation.

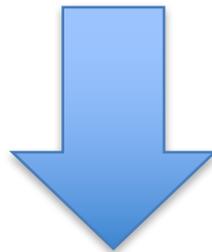
:based_near rdfs:range :Place.

"Unsafe" update:

```
INSERT {?X :based_near?Y} WHERE {?X :locatedIn ?Y .}
```

intrinsic Inconsistencies

can be caught by "safe rewriting"



```
INSERT{?X :based_near ?Y}  
WHERE{?Y :locatedIn ?X .
```

```
MINUS{ {?X1 :locatedIn ?Y}  
UNION {?X :locatedIn ?Y2}}}
```

Copies of the WHERE clause, variables renamed appropriately.

Safe Rewriting – for the SPARQL enthusiasts: MINUS vs. FILTER NOT EXISTS

*Safe rewriting via FILTER NOT EXISTS doesn't work
(Corner case example):*

```
DELETE{?Z a :Place} INSERT{?X :locatedIn ?Y}
WHERE{ {?Y :locatedIn ?X}
        UNION {?V :locatedIn ?Z}}
```

Safe Rewriting – for the SPARQL enthusiasts: MINUS vs. FILTER NOT EXISTS

Safe rewriting via FILTER NOT EXISTS doesn't work:

```
DELETE{?Z a :Place} INSERT{?X :locatedIn ?Y}
WHERE{ {?Y :locatedIn ?X}
       UNION {?V :locatedIn ?Z}}
```

disjoint sets of variables

```
FILTER NOT EXISTS{
  {?X1 :locatedIn?Y} UNION {?V1 :locatedIn?Z1}
}
```

```
FILTER NOT EXISTS{
  {?X :locatedIn?Y2} UNION {?V2 :locatedIn ?Z2}
}
```

Exists whenever the WHERE clause is satisfied!

Simply renaming the whole WHERE clause is not possible.

Safe Rewriting – for the SPARQL enthusiasts: MINUS vs. FILTER NOT EXISTS

Safe rewriting via MINUS: works!

```
DELETE{?Z a :Place} INSERT{?X :locatedIn ?Y}
WHERE{ {?Y :locatedIn ?X}
        UNION {?V :locatedIn ?Z}}
```

```
MINUS{
  {?X1 :locatedIn?Y} UNION {?V1 :locatedIn?Z1}
```

```
}}
```

Extra union branches do not matter!

```
MINUS{
  {?X :locatedIn?Y2} UNION {?V2 :locatedIn?Z2}
```

```
}}
```

- MINUS removes variable bindings of the WHERE clause *that can be combined with some result of the query in its right-hand side.*
- Only variables from the left-hand side of MINUS are "visible" in its right-hand side: great for our case!

So, how can we do SPARQL updates that preserve consistency (and materialization?)

- Dealing with Different forms of inconsistencies:
 - Intrinsic inconsistencies within updates
 - ... solution: "safe" rewriting
 - **Inconsistencies between "old" and "new" knowledge**
 - ... **Different solution strategies:**
 - **Brave**
 - **Cautious**
 - **Fainthearted (somewhere in between ;-)**

SPARQL updates: deal with inconsistency w.r.t. the old knowledge

Idea: adapt Sem_2^{mat} (rewriting-based) semantics



- **Brave:** when in conflict, prefer new knowledge

– cf. FastEvol [Calvanese et al 2010]



- **Cautious:** when in conflict, stick to the old knowledge

– In batch updates, allow variable bindings only where the insert clause does not produce a clash



- **Fainthearted:** relaxation of cautious semantics

– the same batch update might resolve clashes by deleting conflicting parts of the old knowledge!

Example: Brave Sem_2^{mat}

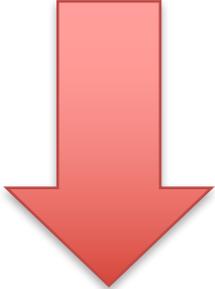
:Place owl:disjointWith :Organisation .

:based_near rdfs:domain :Organisation.

:based_near rdfs:range :Place.



INSERT{?X :based_near?Y} WHERE{?Y :locatedIn ?X}

 ***Preprocess:
safe rewriting***

INSERT{?X :based_near?Y}

WHERE{?Y :locatedIn ?X

MINUS{ {?Y1 :locatedIn?X}

UNION {?Y :locatedIn?X2}}}

Example: Brave Sem_2^{mat}

:Place owl:disjointWith :Organisation .
:based_near rdfs:domain :Organisation.
:based_near rdfs:range :Place.



INSERT{?X : based_near ?Y} WHERE{?X :locatedIn ?Y}



Brave Sem_2^{mat}

DELETE {?X a :Place . ?X3 :based_near ?X .
?Y a :Organization. ?Y :based_near ?Y3 }

Potential
clashes

INSERT{?X :based_near ?Y . ?X a :Organisation . ?Y a :Place}
WHERE{?X :locatedIn ?Y

MINUS{ {?Y1 :locatedIn ?X}
UNION {?Y :locatedIn ?Y2}}

OPTIONAL {?X3 :based_near ?X}
OPTIONAL {?Y :based_near?Y3}}

Bind variables
in DELETE

Example: Cautious Sem_2^{mat}

:Place owl:disjointWith :Organisation .
:based_near rdfs:domain :Organisation.
:based_near rdfs:range :Place.



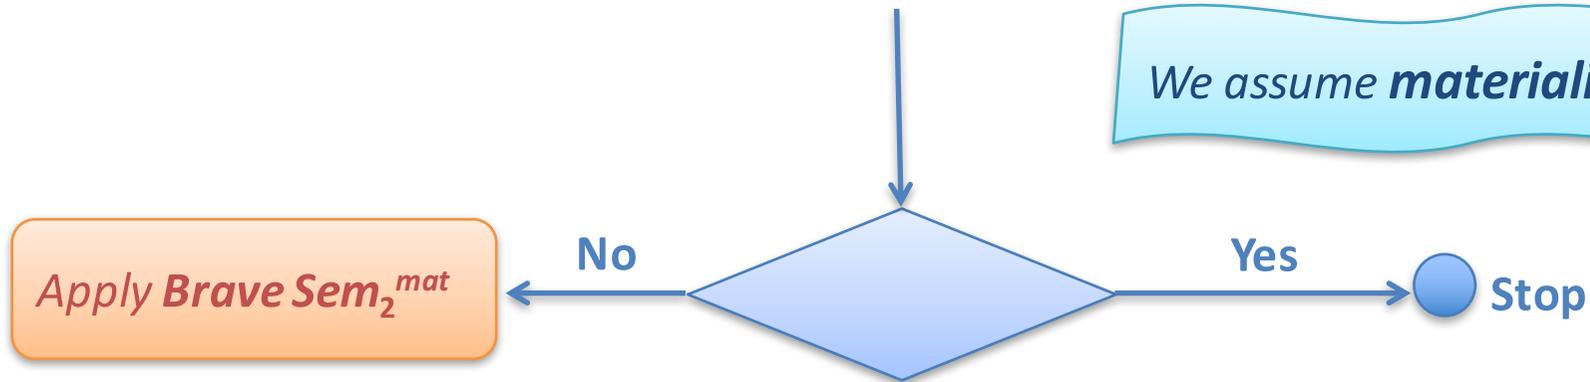
INSERT{?X :based_near ?Y} WHERE{?X :locatedIn ?Y}

↓ **Cautious Sem_2^{mat}**

ASK WHERE{?X :locatedIn ?Y .

{{?X :a :Place} UNION {?Y :a :Organization}}}

We assume *materialised* store!



Example: Cautious Sem_2^{mat}



Removes
some
clashes!

`owl:disjointWith:Organisation`
`near rdfs:domain:Organisation.`
`based_near rdfs:range:Place.`

DELETE{?X :Place}

INSERT{?X :based_near ?Y} WHERE{?X :locatedIn ?Y}

↓ **Cautious Sem_2^{mat}**

ASK WHERE{?X :locatedIn ?Y .

~~{?X :a :Place} UNION {?Y :a :Organization}}~~

Handled by
"DELETE" ... (we
don't want to be
too cautious)

Apply **Brave Sem_2^{mat}**

No

Yes

Stop

Finally:

Example: Fainthearted Sem_2^{mat}

`:Place owl:disjointWith :Organisation .`

`!self domain:Organisation .`

Do inserts only with non-clashing variable bindings



DELETE{?X :Place}

INSERT{?X :based_near ?Y} WHERE{?X :locatedIn ?Y}



Fainthearted Sem_2^{mat}

DELETE{?X :Place . }

INSERT{?X :based_near ?Y . ?X a :Organisation . ?Y a :Place}

WHERE{?X :locatedIn ?Y MINUS{ {?Y1 :locatedIn ?X}
UNION {?Y :locatedIn ?Y2}}

MINUS {{?X a :Place} UNION {?Y a :Organisation}}

Again, handled
by "DELETE"

Fainthearted semantics: pitfalls, e.g. clashes removed by **different** bindings

```
DELETE {?Z a :Place}
INSERT {?X :based_near ?v1}
WHERE { ... }
```

$\mu_2 = [..., ?Z \mapsto :NLC, ...]$: Clears the clash!

$\mu_1 = [?X \mapsto :NLC, ...]$: **clash**

Old state:
:NLC a :Place

- Atomic updates: for each variable binding μ of the WHERE clause either **both** delete **and** insert or **none**.
- Insert with μ_1 depends on the deletion with μ_2 ... our initial approach would be too cautious.
- By atomiticy, if μ_2 also causes insertion (which might depend on the deletion by some μ_3 , etc).

Idea: give up on update atomicity. Delete for all μ_i of the WHERE pattern, insert only where not clashing; for this we have to "separate" DELETE and INSERT... More involved rewriting → paper

Putting the pieces together: What else you find in the paper

- Details, general rewriting algorithms for Brave, Cautious and Feinthearted Update Semantics
- Experiment on some updates with LUBM50 (to show feasibility)
 - no clear winner in terms of performance...
 - optimizations are on our agenda.
- Working **prototype**, in principle pluggable on top of arbitrary SPARQL engines, available at:

<http://dbai.tuwien.ac.at/user/ahmeti/sparqlupdate-inconsistency-resolver/>



What's next?

- SPARQL Update + Entailments
 - from the “*one ring*” to the “*Holy Grail*”:
 - SPARQL Updates for full OBDA? (i.e. incl mappings)



- Initial work to extend our work to updates over DBPedia **including mappings** - forthcoming!
(sneak preview: short paper at AMW2016, next week, Panama)