

# Simple Transaction-Time ESWC 2016

James Anderson <james@dydra.com>

Arto Bendiken <arto@dydra.com>

@dydradata @lomoramic @bendiken

<<http://dydra.com/>>



# Overview

- Temporal RDF processing : alternative ...
  - Data models
  - Processing models
  - Expression
- Exemplified by
  - Archival : Diachron
  - Stream : C-SPARQL, CQELS, SPARQLStream, EP-SPARQL, INSTANS, R&WBase, STRABON
- Dydra : in terms of transaction time
  - capabilities
  - benefits
- Goal : raise three questions

# Three Questions

- How to align levels?
  - interfaces and protocols
  - schematic, conceptual
  - physical, concrete representation
  - is `CLOCK_MONOTONIC` pervasive
- Why monolithic expressions?
  - dataset
  - projection
- Why homogeneous temporality?
  - static, archival
  - dynamic, stream

# Precedent: TSQL

- TSQL<sup>[1][2]</sup>
  - Data :  $(O_+ \times T)$ 
    - concerns the temporal properties of *tuples*
  - Expression

```
CREATE TABLE Employee(  
  Name VARCHAR(30),  
  Manager VARCHAR(30) VALIDTIME REFERENCES Employee (Name), Dept  
  VARCHAR(20)) AS VALIDTIME PERIOD(DATE)
```

```
ALTER TABLE Employee ADD TRANSACTIONTIME
```

```
TRANSACTIONTIME SELECT Dept, COUNT(*) FROM Employee  
GROUP BY Dept  
HAVING COUNT(*) > 25
```

[3]

[1] : <<https://www.cs.arizona.edu/~rts/initiatives/tsql2/tsqldesignapproach.pdf>>

[2] : <<http://timecenter.cs.aau.dk/TimeCenterPublications/TR-1.pdf>>

[3] : <<http://timecenter.cs.aau.dk/TimeCenterPublications/TR-8.pdf>>

# Precedent: Temporal RDF

- Temporal RDF<sub>[1]</sub>
  - Data :  $( (S \times P \times O) \times T )$
  - Processing
  - Expression

*TriplePattern ::= (Subject Predicate Object ) : [T]*

In scenarios where changes are frequent and only affecting a few elements of the document, creating a new physical version of the graph each time an update occurs may lead to large overheads when processing temporal queries that span multiple versions.[1]

[1] : <<http://users.dcc.uchile.cl/~cgutierr/papers/temporalRDF.pdf>>

[2]

# Precedent: R&W Base

- R&W Base<sup>[1]</sup>
  - Data :  $( (S \times P \times O) \times T )$
  - Processing
    - SPARQL rewriting
    - delegation to endpoint
  - Expression

*Datset ::= FROM VersionIri*

[1] : <<http://events.linkedata.org/ldow2013/papers/ldow2013-paper-01.pdf>>

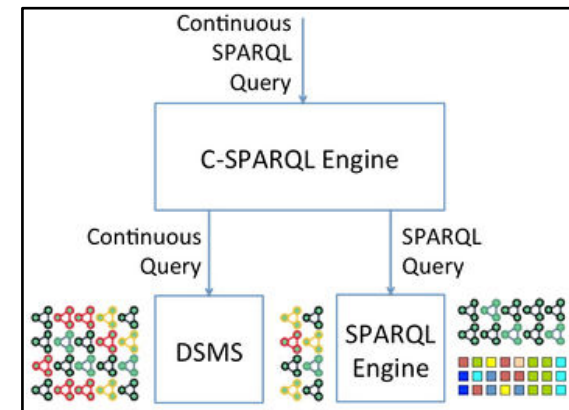
[2]

# Precedent: C-SPARQL

- C-SPARQL<sub>[1]</sub>
  - Data :  $( (S \times P \times O) \times T )$
  - Processing
    - DSMS constitutes snapshot
    - SPARQL processor interprets snapshot as dataset
  - Expression

*Dataset ::= FROM Iri Window ?*

*Window ::= ( RANGE Number TimeUnit ( ( STEP Number TimeUnit ) | TUMBLING ) )  
| ( TRIPLES Number )*



[2]

[1] : <<http://www09.sigmod.org/sigmod/record/issues/1003/p20.article.barbieri.pdf>>

[2] : <<https://www.w3.org/community/rsp/wiki/images/thumb/c/c4/C-sparql-engine.jpg/800px-C-sparql-engine.jpg>>




# Precedent: CQELS

- CQELS<sup>[1]</sup>
  - Data :  $( (S \times P \times O) \times T )^* + ( (S \times P \times O)^* \times T )$
  - Processing
    - Windowing constitutes snapshot dataset
    - Relational SPARQL processor interprets
    - Streaming re-generates a stream
  - Expression

*GraphPatternNotTriples ::= ... | StreamGraphPattern*

*StreamGraphPattern ::= ( RANGE [ Window ] VarOrIRIRReference { TriplesTemplate } )*

*Window ::= ( RANGE Number TimeUnit ( SLIDE Number TimeUnit ) ? )  
| ( TRIPLES Number ) | NOW | ALL*



[1] : <<http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.361.1369>>

[2] : <<https://www.w3.org/community/rsp/wiki/images/7/7f/Cqels-engine.jpg>>

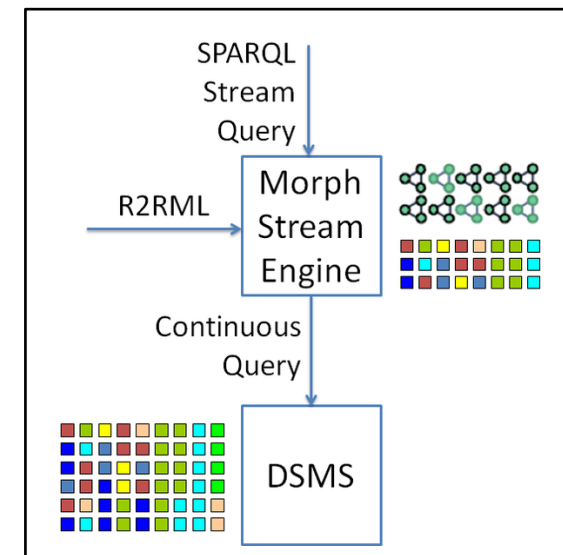


# Precedent: SPARQL Stream

- SPARQL Stream<sup>[1]</sup>
  - Data :  $( (S \times P \times O) \times T )$
  - Processing
    - transform and delegate to DSMS
  - Expression

*Datataset ::= FROM STREAM Iri Window ?*

*Window ::= FROM NOW ( - Number TimeUnit )?  
          TO NOW ( - Number TimeUnit )?  
          ( STEP Number TimeUnit ) ?*



[1] : <[http://oa.upm.es/5121/1/Enabling\\_Ontology-based\\_Access\\_to\\_Streaming\\_Data\\_Sources.pdf](http://oa.upm.es/5121/1/Enabling_Ontology-based_Access_to_Streaming_Data_Sources.pdf)>

[2] : <<https://www.w3.org/community/rsp/wiki/File:Morph-streams.png>>



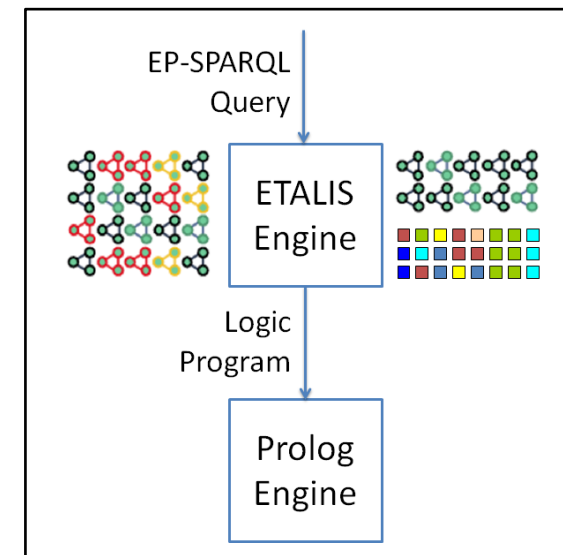
# Precedent: EP-SPARQL

- EP-SPARQL[1]
  - Data :  $(S \times P \times O) \times T \times T$
  - Processing :
    - transform and delegate to Prolog
  - Expression

*Dataset ::=*

*GraphPatternNotTriples ::= ... | SeqGraphPattern*  
| *EqualsGraphPattern*  
| *OptionalSeqGraphPattern*  
| *OptionalEqualsGraphPattern*

*NullOperator ::= ... | getDURATION*  
| *getENDTIME*  
| *getSTARTTIME*



[2]

[1] : <<http://www.wwwconference.org/proceedings/www2011/proceedings/p635.pdf>>

[2] : <<https://www.w3.org/community/rsp/wiki/images/c/cb/Etalis.png>>

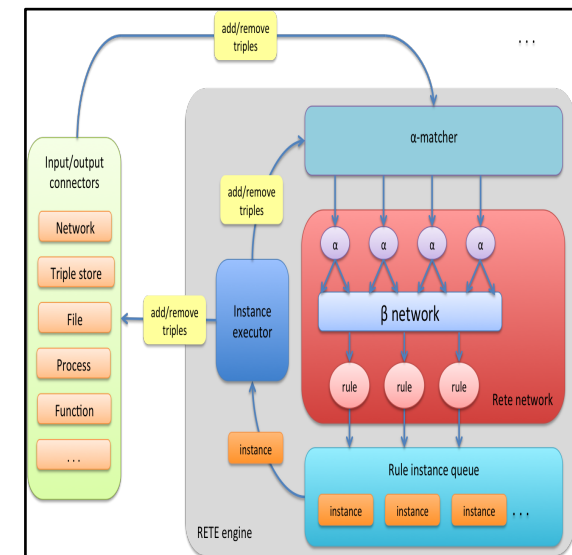


# Precedent: INSTANS

- INSTANS<sub>[1]</sub>
  - Data :  $(S \times P \times O)$
  - Processing :
    - Rete network propagation
  - Expression

*Datatset ::=*

*Verb ::= ... | :hasWindow*



[1] : <[http://iswc2012.semanticweb.org/sites/default/files/paper\\_22.pdf](http://iswc2012.semanticweb.org/sites/default/files/paper_22.pdf)>

[2] : <[http://iswc2012.semanticweb.org/sites/default/files/paper\\_22.pdf](http://iswc2012.semanticweb.org/sites/default/files/paper_22.pdf)>

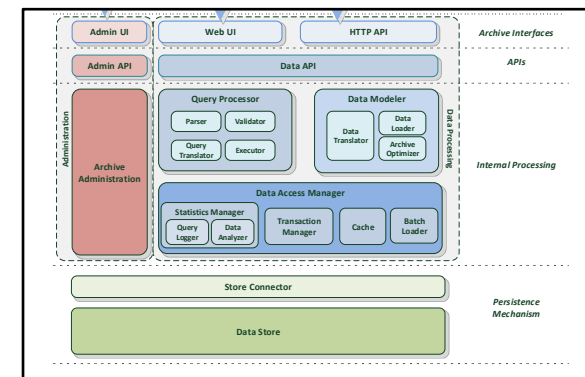


# Precedent: DIACHRON

- Diachron<sub>[1]</sub>
  - Data :  $(S \times P \times O)$
  - Processing :
    - SPARQL evaluation of DIACHRON query
  - Expression

```
SourceClause ::= ... | ( FROM DATASET URI ( AT VERSION URI )? )
                | ( FROM CHANGES URI ( BEFORE VERSION URI )? )
                | ( AFTER VERSION URI )
                | ( BETWEEN VERSIONS URI URI )
```

```
SourcePattern ::= ... | ( DATASET URI ( AT VERSION URI )? )
                 | ( CHANGES URI ( BEFORE VERSION URI )? )
                 | ( AFTER VERSION URI )
                 | ( BETWEEN VERSIONS URI URI )
```



[2]

[1] : <<http://arxiv.org/pdf/1504.01891.pdf>>

[2] : <<http://arxiv.org/pdf/1504.01891.pdf>>



# Alternatively: Dydra

- Dydra<sup>[1]</sup>
  - Data :  $( (S \times P \times O \times G)^* \times R )$
  - Processing :
    - SPARQL, semantics analogous to named graphs
  - Expression:

*GraphPatternNotTriples ::= ... | RevisionGraphPattern*

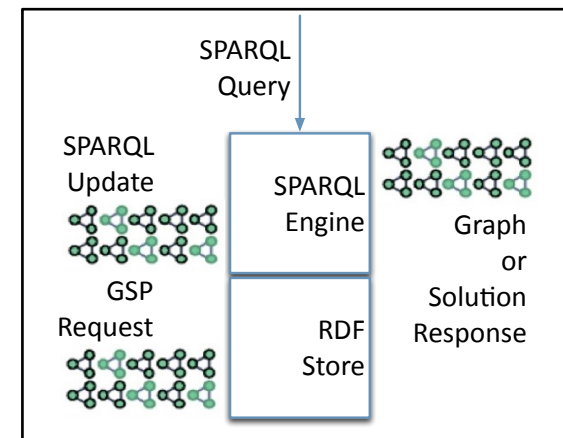
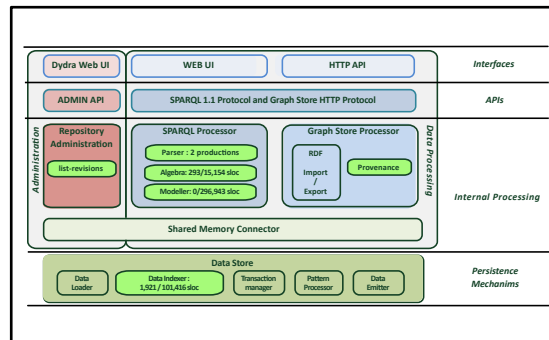
*RevisionGraphPattern ::= REVISION ( Var | RevisionRef | String ) GroupGraphPattern*

*RevisionRef ::= ( R Integer ? / ) ? RevisionRelation [3]*

*RevisionRelation ::= ( Revision ( / Revision ) ? ) | ( / Revision ) | XPathDuration )  
( / ( XPathDuration | Integer ) ) ?*

*Revision ::= UUID | /HEAD(~[0-9]+)?/ | XPathDateTime*

*NullOperator ::= ... | THEN*



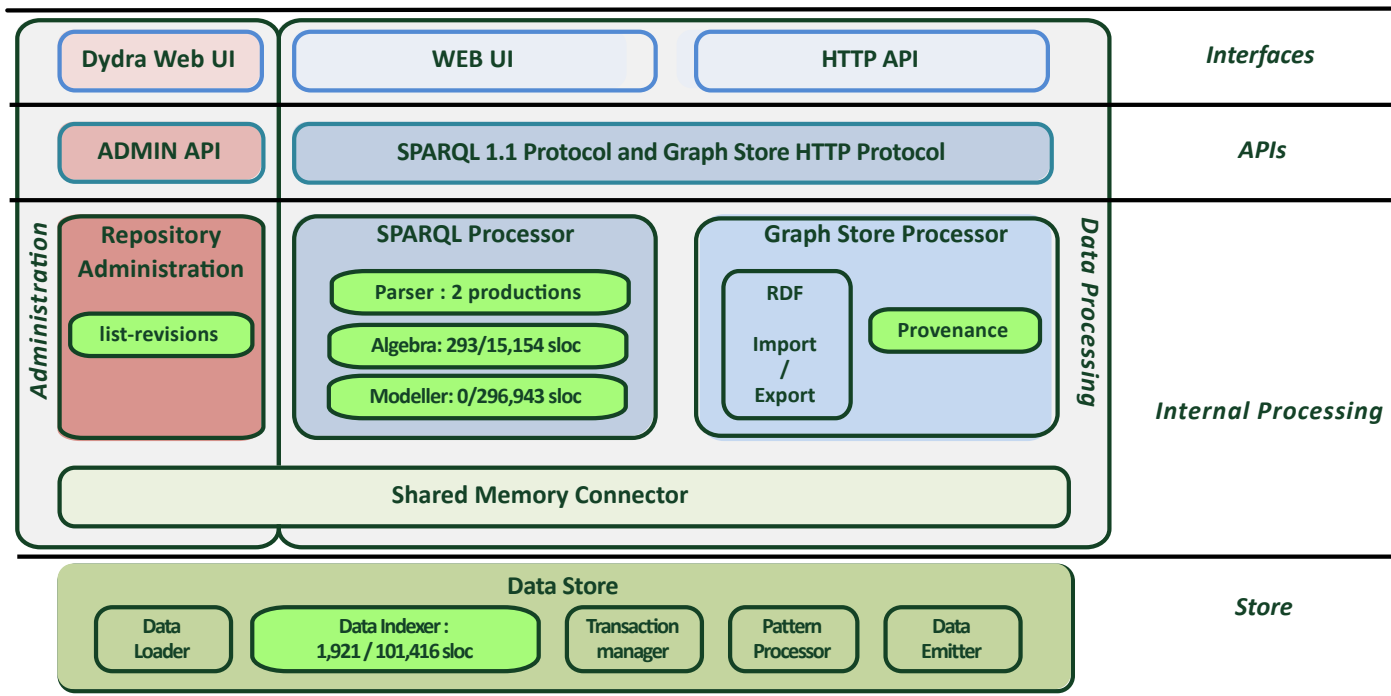
[1] : <<http://dydra.com>>

[2] : <<http://arxiv.org/pdf/1504.01891.pdf>> (sampled)

[3] : <[https://en.wikipedia.org/wiki/ISO\\_8601](https://en.wikipedia.org/wiki/ISO_8601)>



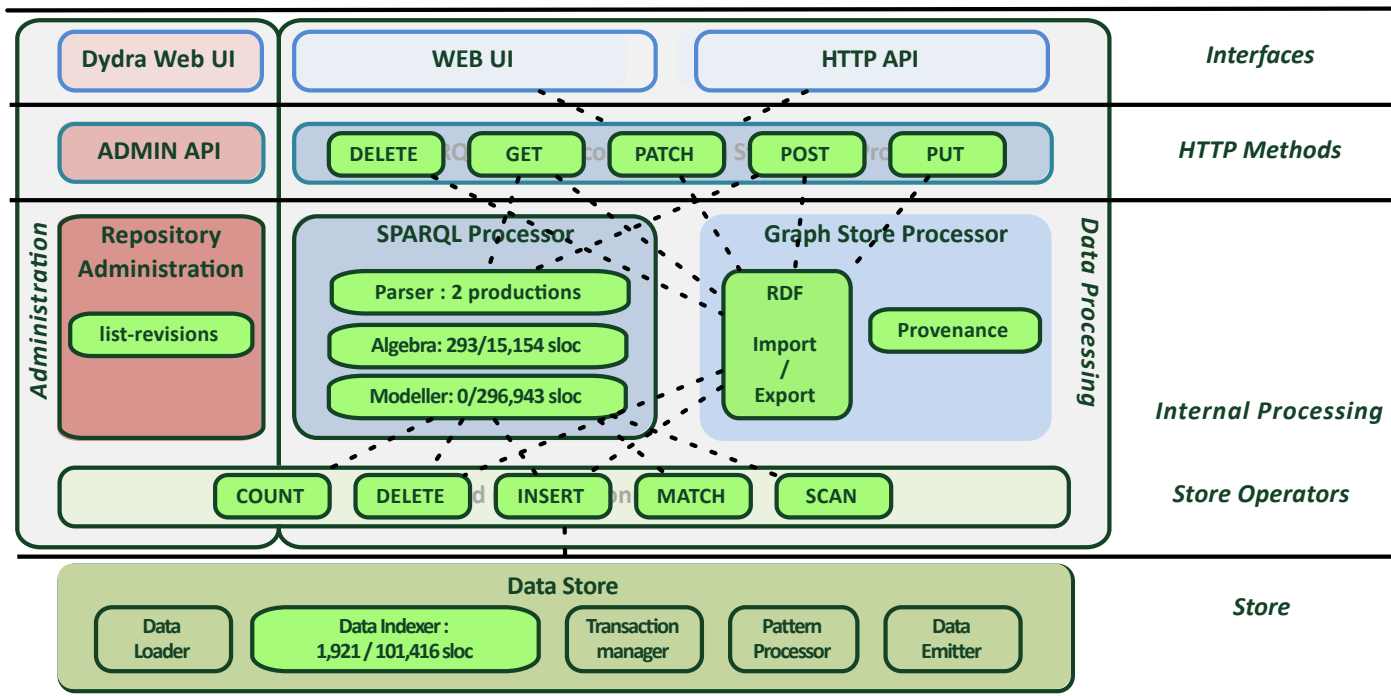
# Dydra: Architecture



[1] : <<http://dydra.com>>

[2] : ◇ sampled

# Dydra: Data-Flow



[1] : <<http://dydra.com>>

[2] : ◇ sampled

# By Analogy to Graph Provenance

- Named graphs originated as provenance designators
- Evolved to permit assertions about statement sets
- Temporal coincidence is analogous to physical or conceptual
  - but neither equivalent nor autonomous
- Add a REVISION clause to the SPARQL -- RDF API

*RevisionGraphPattern ::= REVISION ( Var | RevisionRef | String ) GroupGraphPattern*

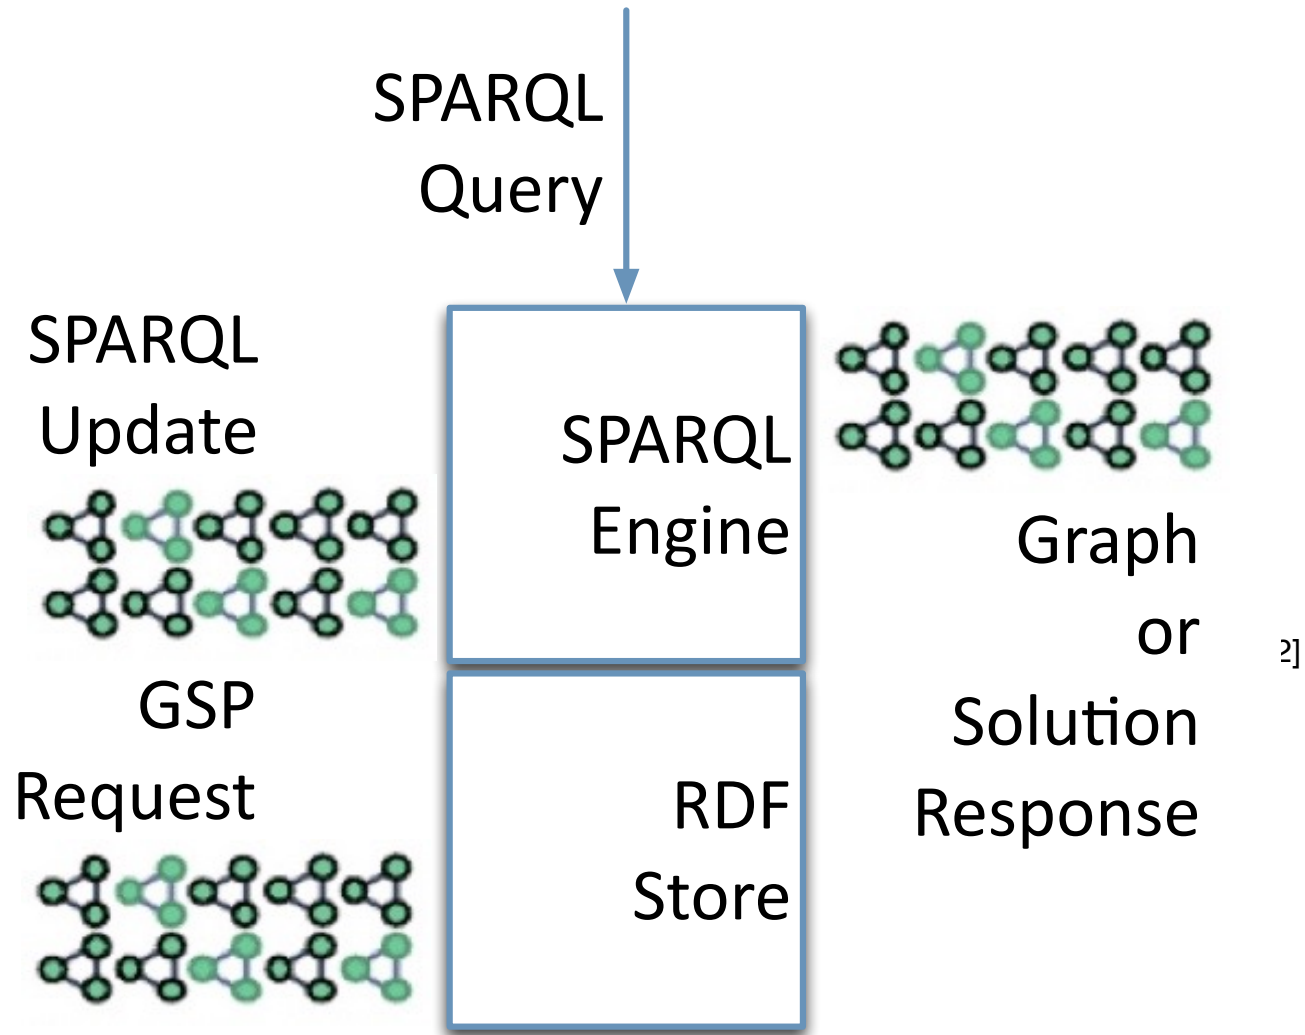
# Dydra : Revision Designators

- Elementary
  - absolute: *UUID* : 01234567-0123-0123-012345678901
  - relative : */HEAD(~[0-9]+)?/* : HEAD~1, HEAD~1..HEAD
  - temporal : *XPathDateTime* : 20150303T180000Z
- Compound
  - interval : *ISOInterval* : 20150303T180000Z/PT10M
  - stream : *ISORepeatingInterval* : R/20150303T180000Z/PT10M/PT01M

# Dydra : Revision Clauses

- Constant
  - **absolute** : `REVISION UUID { ... }`
  - **relative** : `REVISION /HEAD(~[0-9]+)?/ { ... }`
  - **temporal** : `REVISION XPathDateTime { ... }`
- Variable
  - **SIP** : `{ ?subject :revision ?revision } ... REVISION ?revision { ... }`
  - **Protocol** : `HTTP://...?REVISION=UUID ... REVISION ?revision { ... }`
- **Intra-Repository** : `REVISION ?revision { ?subject a ?class }`
- **Extra-Repository** : `SERVICE ?revisionEndpoint { ?subject a ?class }`

# Dydra: RSP



```
PREFIX c: <http://linkedurbandata.org/city#>
PREFIX t: <http://linkedurbandata.org/traffic#>
```

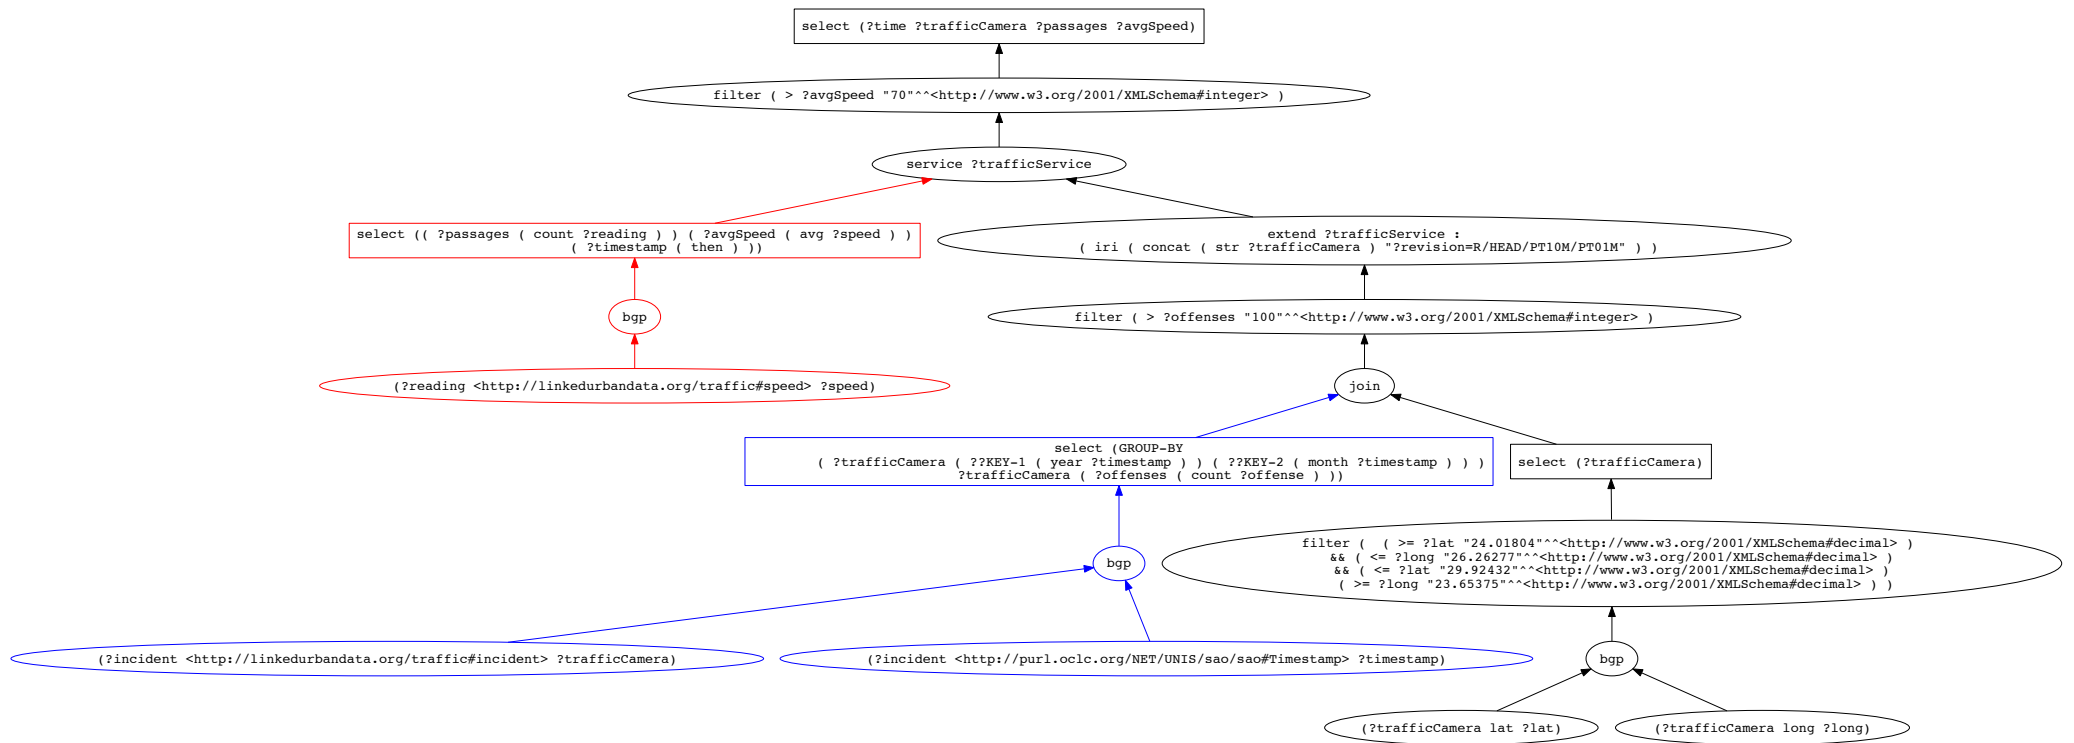
```
select ?time ?trafficCamera ?passages ?avgSpeed where {
# select tollgates in some region (static base data)
{{ select ?trafficCamera where {
  ?trafficCamera <http://www.w3.org/2003/01/geo/wgs84_pos#lat> ?lat .
  ?trafficCamera <http://www.w3.org/2003/01/geo/wgs84_pos#long> ?long .
  filter ( (?lat >= 24.01804) && (?lat <= 29.92432)
    && (?long >= 23.65375) && (?long <= 26.26277) )
} }
# reduce that set to those with some activity level
{ select ?trafficCamera (count(?offense) as ?offenses)
  where {
    ?incident t:incident ?trafficCamera .
    ?incident <http://purl.oclc.org/NET/UNIS/sao/sao#Timestamp> ?
timestamp .
  } group by ?trafficCamera (year(?timestamp)) (month(?timestamp))
}
  filter (?offenses > 100) }
# look for incidents at that location
bind (iri(concat(str(?trafficCamera), '?revision=R/HEAD/PT10M/PT01M'))
as ?trafficService)
service ?trafficService {
  select (count(?reading) as ?passages)
    (avg(?speed) as ?avgSpeed)
    (then() as ?timestamp)
  where { ?reading t:speed ?speed .}
}
  filter (?avgSpeed > 70)
}
```

```

(select
  (service ?trafficService
    (select
      (bgp (triple ?reading <http://linkedurbandata.org/traffic#speed> ?speed))
      (:HAVING (> ?avgSpeed 70) (?passages (count ?reading))
        (?avgSpeed (avg ?speed)) (?timestamp (then))))))
  (extend
    (join
      (select
        (bgp
          (triple ?incident <http://linkedurbandata.org/traffic#incident> ?trafficCamera)
          (triple ?incident <http://purl.oclc.org/NET/UNIS/sao/sao#Timestamp> ?timestamp))
        (:HAVING (> ?offenses 100) :GROUP-BY
          (?trafficCamera (??KEY-1 (year ?timestamp)) (??KEY-2 (month ?timestamp)))
          ?trafficCamera (?offenses (count ?offense))))))
      (select
        (filter
          (bgp (triple ?trafficCamera <http://www.w3.org/2003/01/geo/wgs84_pos#lat> ?lat)
            (triple ?trafficCamera <http://www.w3.org/2003/01/geo/wgs84_pos#long> ?long))
          (&& (>= ?lat 30623/1275)
            (&& (<= ?long 39079/1488)
              (&& (<= ?lat 35191/1176) (>= ?long 18923/800))))))
        (?trafficCamera)))
    ?trafficService
    (iri
      (concat (str ?trafficCamera) "?revision=R/HEAD/PT10M/PT01M")))
    :QUERY-TEXT
    "SELECT ( COUNT ( ?reading ) AS ?passages ) ( AVG ( ?speed ) AS ?avgSpeed ) ( THEN ( )
  AS ?timestamp ) WHERE { ?reading <http://linkedurbandata.org/traffic#speed> ?speed . } HAVING ( ?
  avgSpeed > 70 ) "
    (?trafficCamera ?time ?passages ?avgSpeed))

```

# Dydra : Processing Model



# Dydra: Examples

- <http://dydra.com/schema-org-test/schema-org>
  - <http://dydra.com/schema-org-test/provenance>
- <http://dydra.com/schema/efo>
  - <http://dydra.com/schema/efo-provenance>

# Dydra: Query @Revision

The screenshot shows the Dydra web interface. At the top, there is a green header with the Dydra logo and navigation links: Home, About, Docs, Blog, My Account, and Logout. Below the header, a dark bar displays the current repository path 'schema-org-test / schema-org' and buttons for 'Query Log', 'Endpoint URL', and 'Back to Repository'. The main content area is divided into a left sidebar and a central editor. The sidebar has a 'Views' section with a list of views: 'Count of properties', 'Count of types', 'Deprecations', 'Medical properties', 'Medical schemas', and 'revisions' (which is selected). Below the views is a 'Default Prefixes' section with a question mark icon. The central editor has a 'Name:' field containing 'revisions' and buttons for 'Run', 'Save', 'Save as', and 'Clear'. The editor contains a SPARQL query:


```
1 select ?revision ?date
2 where {
3   revision ?revision {
4     bind (then() as ?date)
5   }
6 }
7
8
```

At the bottom of the interface, there is a footer with the text 'A product of Datagraph, Inc.' and navigation links: Home, About, Docs, Blog, My Account, Logout, Legal, Support, GitHub, Twitter, and RSS.



# Dydra: View @Revision

dydra.com/schema-org-test/schema-org/revisions.html


 Download ▾

revision	date
<a href="#">urn:uuid:ad278158-020b-b940-a137-f721b69fa314</a>	2015-09-01T10:01:02.771288Z
<a href="#">urn:uuid:604759b6-4a75-b547-b1e1-29589cca6948</a>	2015-09-01T10:00:51.96827Z
<a href="#">urn:uuid:0e65df72-ae2e-bc42-81af-ef8d6756f97d</a>	2015-09-01T10:00:40.918747Z
<a href="#">urn:uuid:58acf71d-04e2-dd4a-98cb-081a46d809a2</a>	2015-09-01T10:00:30.002933Z
<a href="#">urn:uuid:a13c21b6-4c14-db42-85a4-bf5dd34db0e1</a>	2015-09-01T10:00:18.995734Z
<a href="#">urn:uuid:4cfa489e-c6b8-9b4c-96e2-bd0fd48d4f50</a>	2015-09-01T10:00:08.27253Z
<a href="#">urn:uuid:879ccfcc-52c2-8049-8d43-23772758b8c4</a>	2015-09-01T09:59:57.491989Z
<a href="#">urn:uuid:133ac8b1-83b7-ef43-a8ea-c948511402ac</a>	2015-09-01T09:59:46.670262Z
<a href="#">urn:uuid:50bd2160-0069-7a49-80a2-174a76f9576b</a>	2015-09-01T09:59:35.789613Z
<a href="#">urn:uuid:3be0878b-61a7-c846-91de-433e69937dae</a>	2015-09-01T09:59:24.893212Z
<a href="#">urn:uuid:f57db3d0-68aa-f44f-b9e2-4d822e31529a</a>	2015-09-01T09:59:14.034819Z
<a href="#">urn:uuid:b51af61c-1568-944e-a0c6-1be3e860a469</a>	2015-09-01T09:59:03.213645Z


# Dydra: Query @Revision

```
select *
where {
  { revision <urn:uuid:604759b6-4a75-b547-b1e1-29589cca6948> {
    { select (count(*) as ?count) where {?s ?p ?o} }
    bind (then() as ?date)
    bind ("as UUID" as ?form)
  }
} union
{ revision "HEAD~1" {
  { select (count(*) as ?count) where {?s ?p ?o} }
  bind (then() as ?date)
  bind ("relative" as ?form)
}
} union
{ revision "2015-09-01T09:57:07" {
  { select (count(*) as ?count) where {?s ?p ?o} }
  bind (then() as ?date)
  bind ("by date" as ?form)
}
}
}
```

# Dydra: View @Revision



## specific revisions

 Download ▾

count	↕	date	↕	form	↕
11094		2015-09-01T10:00:51.96827Z		as UUID	
11094		2015-09-01T10:00:51.96827Z		relative	
2771		2015-09-01T09:57:06.177342Z		by date	





# Dydra: as SPARQL endpoint

```
select ?versionId ?title ?term ?comment
where {
  { service <http://localhost/schema-org-test/provenance> {
    select ?location ?versionId ?endpoint ?title
    where {
      ?revision <http://www.w3.org/1999/02/22-rdf-syntax-ns#type> <urn:dydra:Revision> .
      { graph ?revision {
        ?revision <http://www.w3.org/ns/sparql-service-description#endpoint> ?endpoint .
        ?revision <http://purl.org/dc/elements/1.1/title> ?title .
        ?revision <http://usefulinc.com/ns/doap#revision> ?versionId .
      }
    }
  }
}
service ?endpoint {
  ?term <http://www.w3.org/2000/01/rdf-schema#comment> ?comment .
  filter (regex(?comment, '.*deprecat.*'))
}
}
```

# Storage Model: Evolution

1. Next, "change based" @statement id
  - efficient HEAD access
  - **inefficient urn:uuid? access**
  - space efficient
  - **inefficient mutation at scale**
2. Initially, "independent copies" @quad term id
  - efficient HEAD access
  - efficient urn:uuid? access
  - **space inefficient**
  - **inefficient mutation at scale**
3. Current, "time-slice based" (various)
  - OIΔI HEAD access
  - OIΔI urn:uuid? access
  - space efficient
  - efficient mutation at scale

# Storage Model: Current

- Current, "time-slice based" (various)
  - OIΔI HEAD access
  - OIΔI urn:uuid? access
  - space efficient
  - efficient mutation at scale
- Analogous to T-SPARQL<sub>[1]</sub>, but where  $T$  is anisotropic, with one dimension - transaction time, comprising UUID values and the others as application concerns.

[1] : Grandi : <<http://www-db.deis.unibo.it/~fgrandi/papers/GraphQ10.pdf>>



# Results: Update Processing

- Time : efficient mutation at scale
  - effort proportional to mutation size (+ / -)
- Space
  - initial space requirements (initially sub-optimal)
  - subsequent compaction (2x)
- Detrimental practices
  - zero-sum non-monotonic updates
  - blank nodes


# Query Processing Performance

- indexed revisions v/s materialized
- spurious b+tree

# Space

ontology	revisions	files (MB)	quads		store (MB)	bytes/quad
			total	HEAD		
gist	16	2.9	47,273	849	38.4	851
schema.org	23	19.1	156,150	9,144	40.6	273
STW	7	94.2	771,375	108,967	240.0	311
efo	144	1,384.0	12,807,240	260,503	4,270	358
efo	@2.69	2.8	260,503	260,503	84.6	324

# Thank you



# DYDRA

James Anderson <james@dydra.com>

Arto Bendiken <arto@dydra.com>

@dydradata @lomoramic @bendiken

<<http://dydra.com/>>