Distributed RDF Query Processing and Reasoning in Peer-to-Peer Networks

Zoi KaoudiPostdoctoral researcher

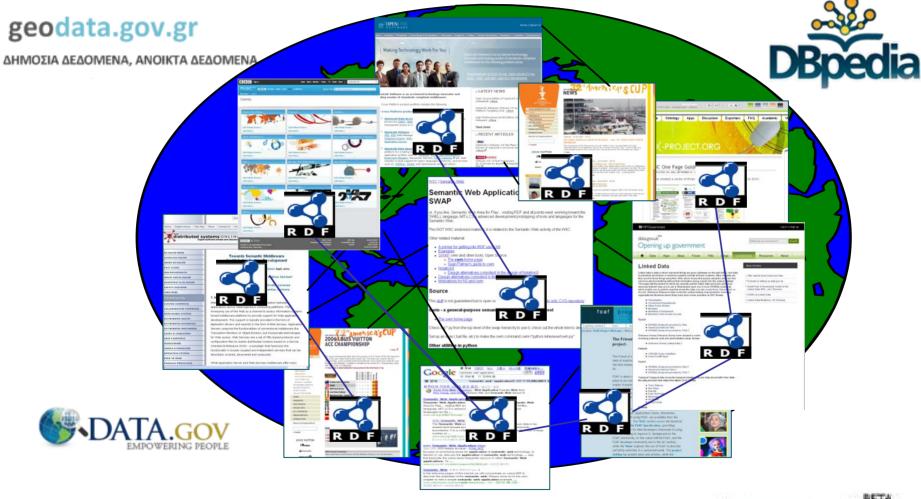




Outline

- Introduction
- The system Atlas
- RDFS Reasoning in P2P networks
- SPARQL Query Processing and Optimization
- Conclusions

Web of Documents/Web of Data

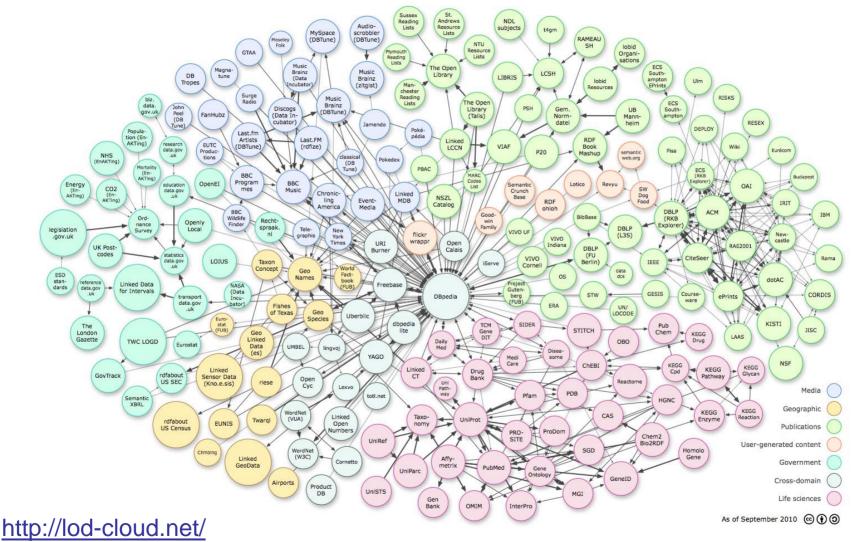


Resource Description Framework



- RDF provides a simple and abstract knowledge representation for resources on the Web which are uniquely identified by *Universal Resource Identifiers* (URIs)
- RDF Schema (RDFS) gives meaning to RDF terms and groups them to concepts
- SPARQL: official W3C recommendation language for querying RDF data

Linked Open Data Cloud



RDF Data Management



 With the vast amount of available RDF data sources on the Web increasing rapidly, there is an urgent need for RDF data management

 RDF storage, query processing and reasoning have been at the center of attention during the last years

RDF Data Management



- Centralized solutions
 - Jena, Sesame, RSSDB, Oracle, RDF-3X, etc.
- Parallel and distributed solutions
 - ← P2P systems
 - RDFPeers, BabelPeers, GridVine, Edutella etc.
 - Powerful clusters
 - Virtuoso, YARS2, MaRVIN
 - Cloud computing
 - Mika08, Urbani09, Stein10

Peer-to-peer (P2P) systems

 P2P systems provide nice features for Internet-scale applications (e.g., content sharing, distributed digital libraries)

- Distributed Hash Tables (DHTs)
 - answer exact match queries efficiently

Challenges

- How to index RDF data and RDFS ontologies in a DHT?
 - Triple indexing scheme and handle RDF and RDFS uniformly
 - > Distributed mapping dictionary
- How to answer SPARQL queries efficiently? What kinds of query optimization techniques to use?
 - Greedy optimization algorithms to minimize the size of the intermediate results
 - Selectivity estimation techniques
 - > RDF statistics in a DHT system
- How to provide reasoning mechanisms for RDFS in a DHT?
 - Distributed forward chaining (FC)
 - Distributed backward chaining (BC)
 - Distributed magic sets transformation (MS)
 - Comparative study (analytically and experimentally)
 - Theoretical proofs of correctness

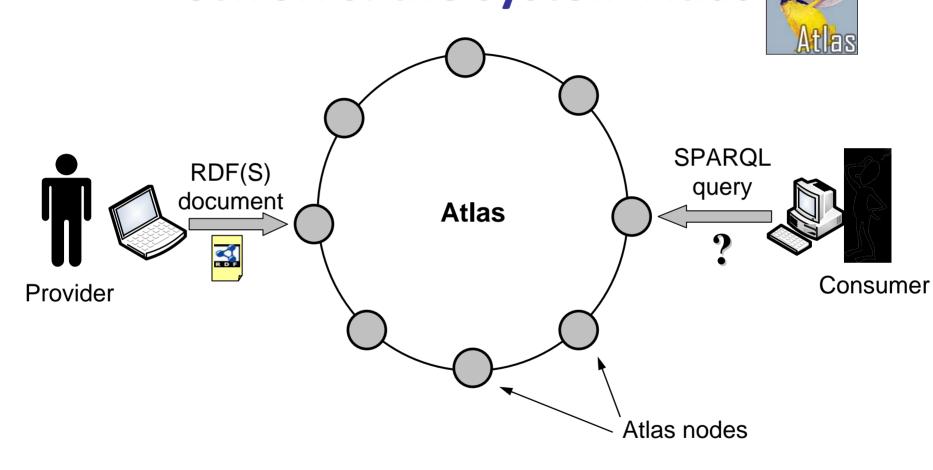
Outline

Introduction

- The system Atlas
- RDFS Reasoning in P2P networks

- SPARQL Query Processing and Optimization
- Conclusions

RDF Data Management in P2P networks: the system Atlas

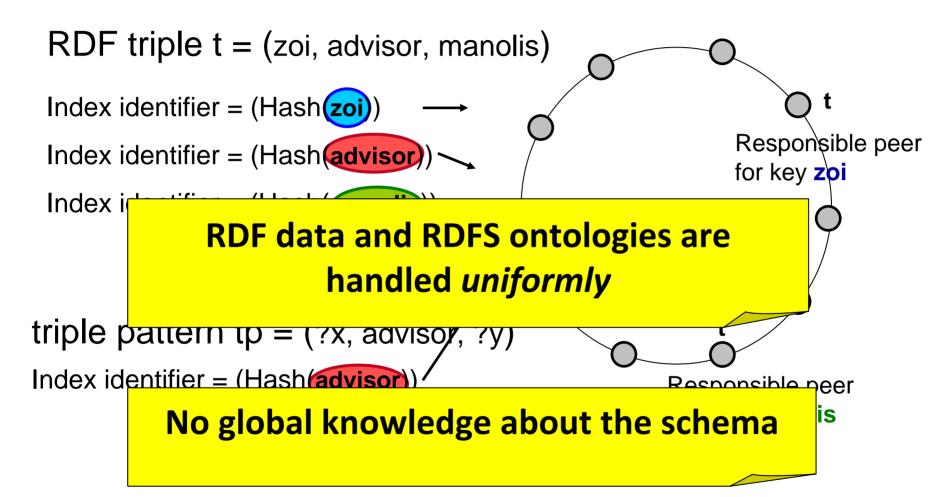


The system Atlas

 Atlas: a full-blown open source P2P system for the distributed processing of RDF and RDFS data stored on top of DHTs (http://atlas.di.uoa.gr)

 Atlas has been used in EU projects OntoGrid and SemsorGrid4Env as a distributed registry of metadata

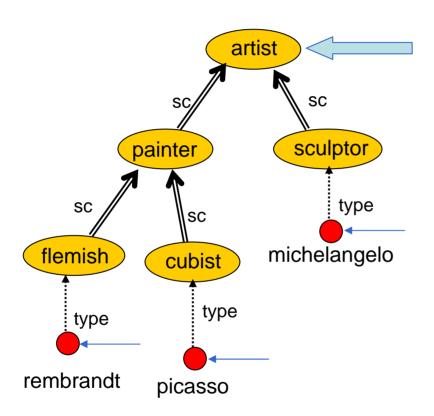
Indexing in Atlas



Outline

- Introduction
- The system Atlas
- RDFS Reasoning in P2P networks
- SPARQL Query Processing and Optimization
- Conclusions

RDFS Reasoning



Query: Find all artists q = (?x, rdf:type, artist)

RDFS entailment rules

?x
rembrandt
picasso
michelangelo

RDFS reasoning techniques

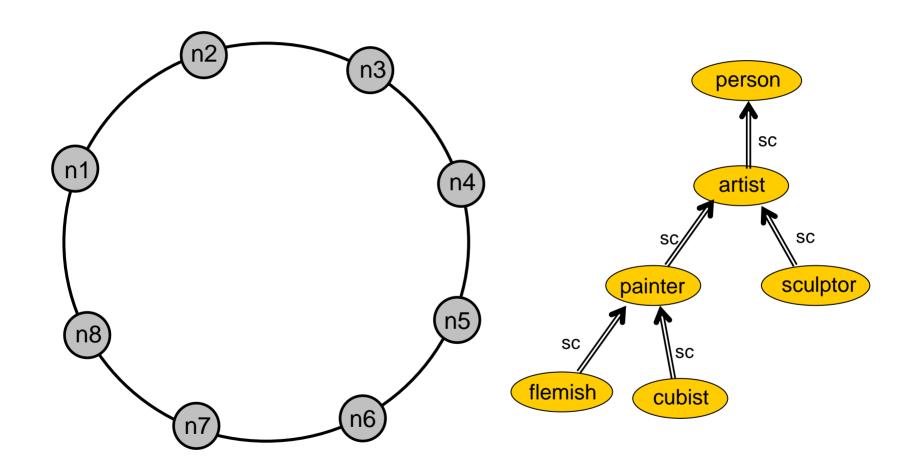
- Bottom-up approach
 - compute all inferences a priori (RDFS closure)
- Top-down approach
 - compute inferences at query run time
- Optimized bottom-up approach
 - compute inferences a priori given a specific query

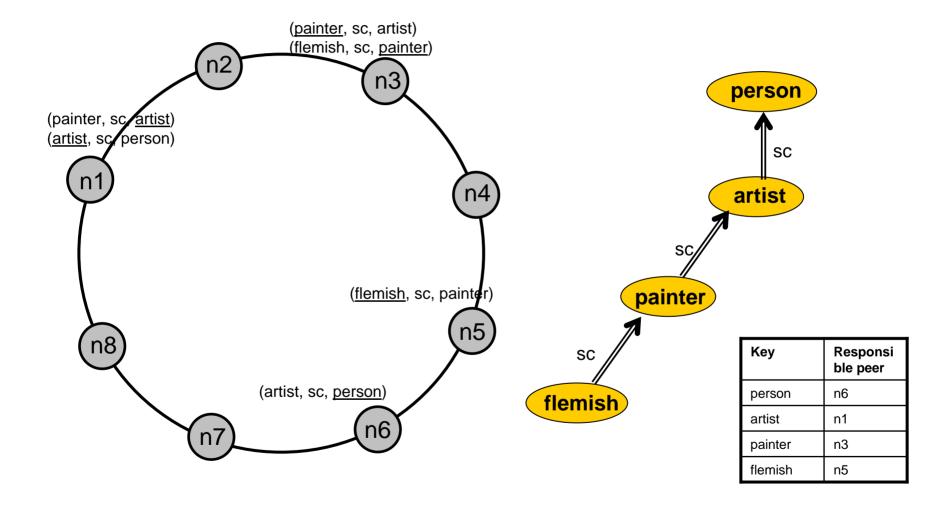
Outline

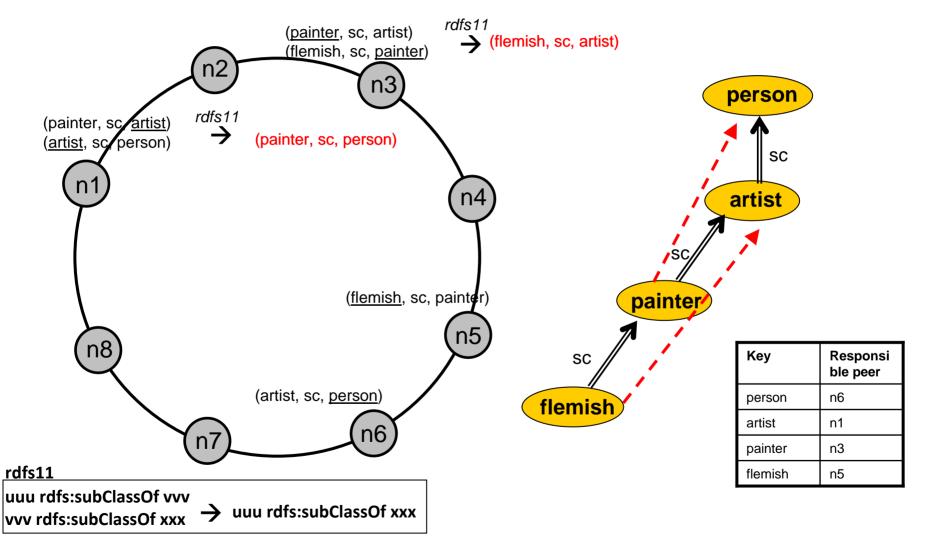
- Introduction
- The system Atlas
- RDFS Reasoning in P2P networks
 - Bottom-up approach
 - Top-down approach
 - Optimized bottom-up approach
 - Evaluation
- SPARQL Query Processing and Optimization
- Conclusions

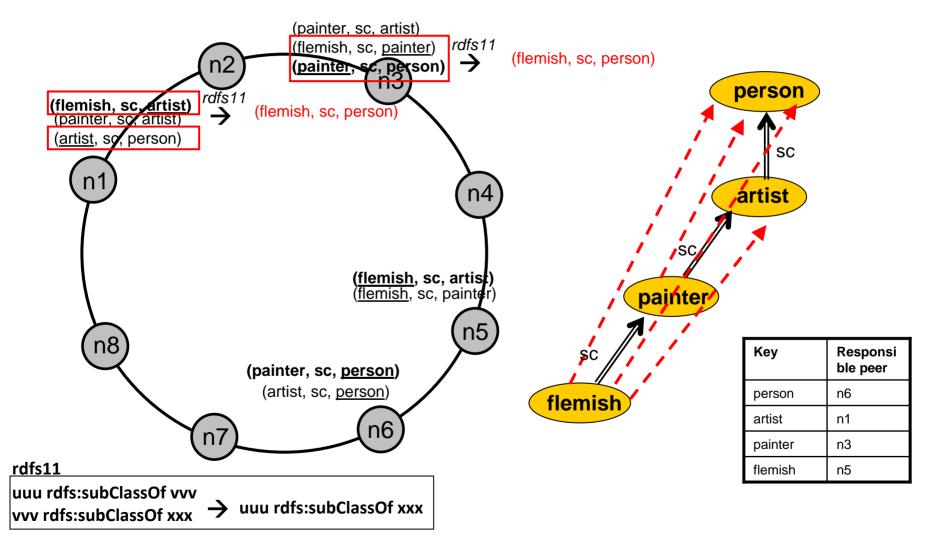
Bottom-up approach

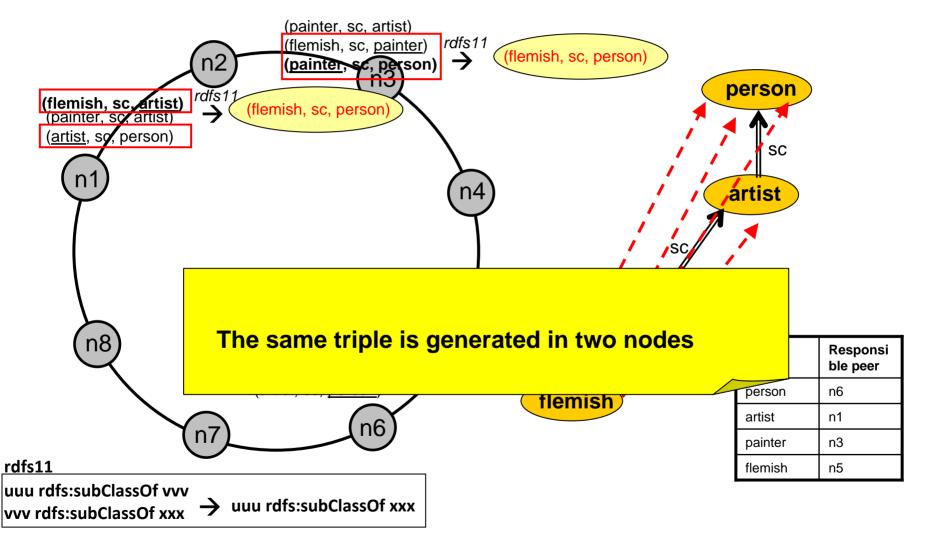
- Distributed forward chaining algorithm
 - compute all inferences a priori (RDFS closure)
 - many redundant triples

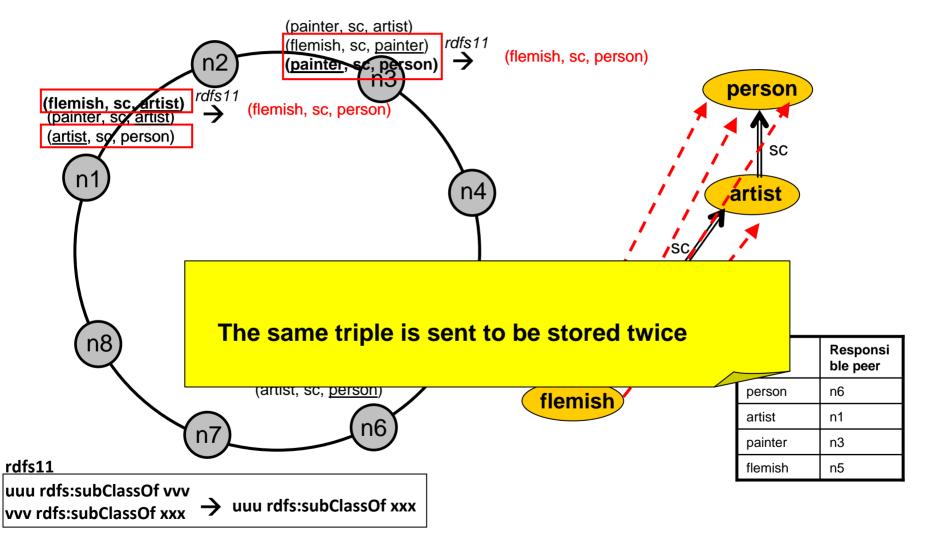












RDFS Entailment Rules

(Only the RDFS rules proposed in the minimal deductive system **mrdf** of [Munoz et al.2009])

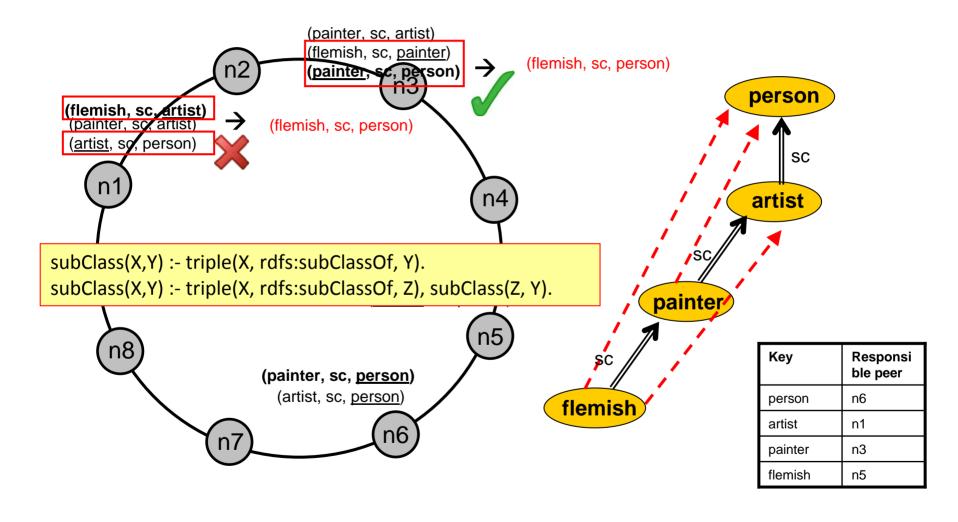
- subClass(X,Y) :- triple(X, rdfs:subClassOf, Y).
- subClass(X,Y) :- triple(X, rdfs:subClassOf, Z), subClass(Z, Y).

edb relation: triple

- subProperty(X, Y) :- triple(X, rdfs:subPropertyOf, Y).
- subProperty(X, Y):- triple(X, rdfs:subPropertyOf, Z), subProperty(Z, Y).

idb relations: subClass, subProperty, newTriple, type

- newTriple(X, P, Y):- triple(X, P, Y).
- newTriple(X, P, Y) :- triple(X, P1, Z), subProperty(P1, P).
- type Safe , Y)
 - type Z, rdfs:subC
- Linear
- type(X, Y) :- triple(X, P, Z), triple(P, rdfs:domain, Y).
- type(X, Y) :- triple(Z, P, X), triple(P, rdfs:range, Y).

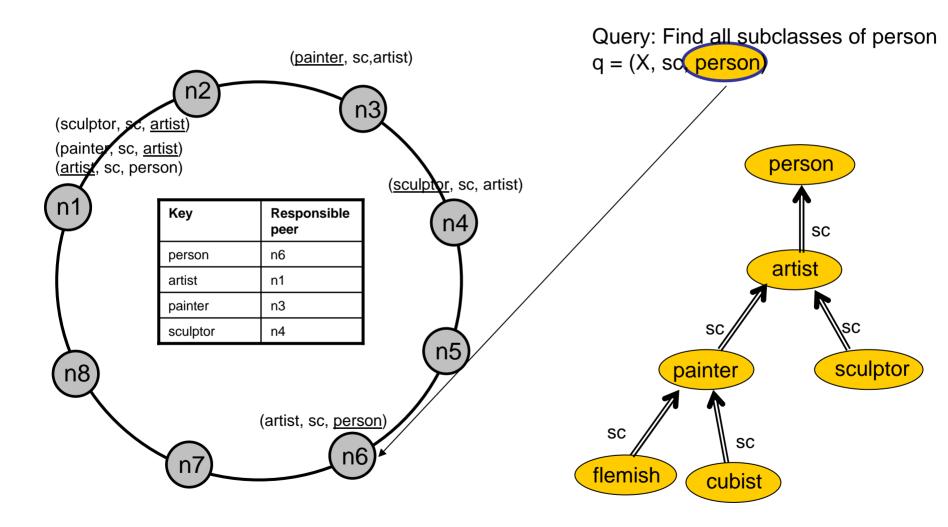


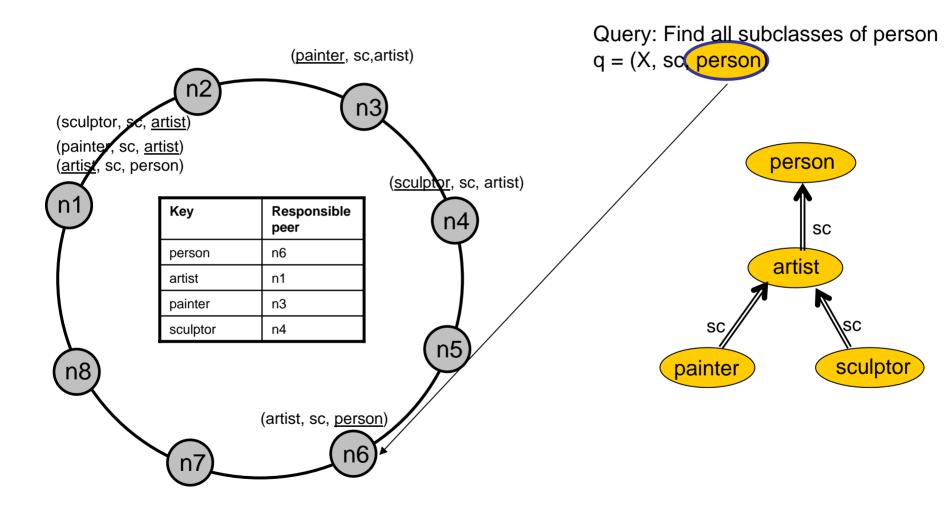
Outline

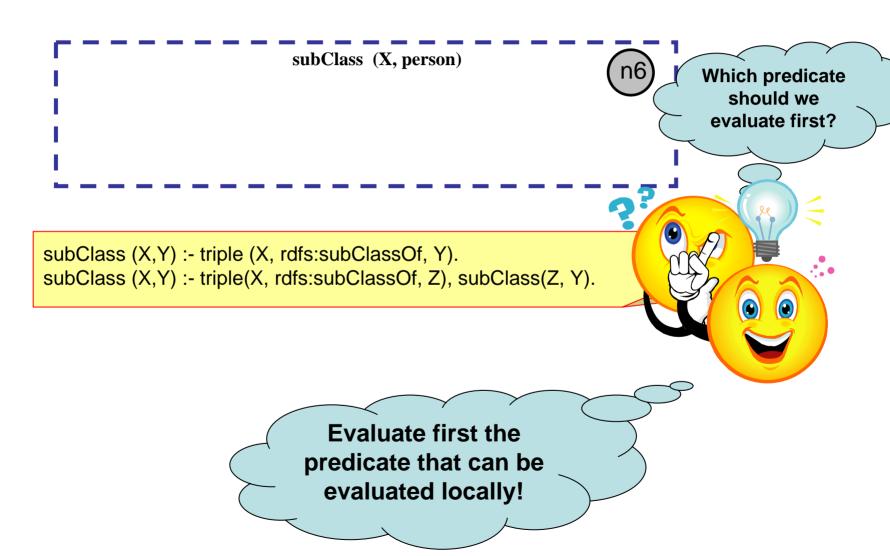
- Introduction
- The system Atlas
- RDFS Reasoning in P2P networks
 - Bottom-up approach
 - Top-down approach
 - Optimized bottom-up approach
 - Evaluation
- SPARQL Query Processing and Optimization
- Conclusions

Top-down approach

- Distributed backward chaining algorithm
 - store only the given triples
 - compute necessary inferences at query run time





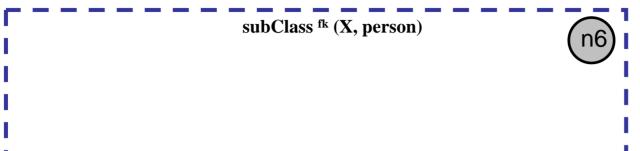


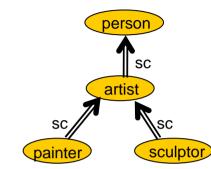
RDFS Entailment Rules - revisited

- Recursive rules
- Rule adornment from recursive query processing
 - Good orderings for evaluating predicates
 - eg. subClass(X, artist) → subClass^{fb} (X,Y)
- Extended adornment
 - Ordered string of f, b, k
 - k: an argument that is bound and the key
 - b : bound argument (not the key)
 - *f* : free argument
 - eg. At node responsible for key artist:
 triple(X, rdf:type, artist) triple^{fbk} (X, rdf:type, Y).
 - Good ordering for evaluating predicates in a distributed environment

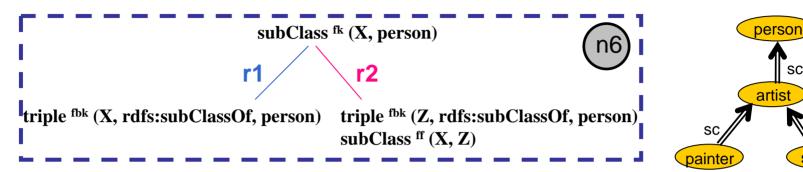
RDFS Entailment Rules - revisited

- subClass^{kf} (X,Y) :- triple^{kbf} (X, rdfs:subClassOf, Y).
- subClass^{kf} (X,Y):- triple^{kbf} (X, rdfs:subClassOf, Z), subClass^{ff} (Z, Y).
- subClass^{fk} (X,Y) :- triple^{fbk} (X, rdfs:subClassOf, Y).
- subClass^{fk} (X,Y) :- subClass^{ff}(X, Z), triple^{fbk} (Z, rdfs:subClassOf, Y).
- subProperty^{kf} (X,Y) :- triple^{kbf} (X, rdfs:subPropertyOf, Y).
- subProperty^{kf} (X,Y) :- triple^{kbf} (X, rdfs:subPropertyOf, Z), subProperty^{ff} (Z, Y).
- subProperty^{fk} (X,Y) :- triple^{fbk} (X, rdfs:subPropertyOf, Y).
- subProperty^{fk} (X,Y):- subProperty^{ff}(X, Z), triple^{fbk} (Z, rdfs:subPropertyOf, Y).
- type^{kf} (X, Y) :- triple^{kbf} (X, rdf:type, Y).
- type^{kf} (X, Y) :- triple^{kff} (X, P, Z), triple^{fbf} (P, rdfs:domain, Y).
- type^{kf} (X, Y) :- triple^{ffk} (Z, P, X), triple^{fbf} (P, rdfs:range, Y).
- type^{kf} (X, Y) :- triple^{kbf} (X, rdf:type, Z), subClass^{ff} (Z, Y).
- type^{fk} (X, Y) :- triple^{fbk} (X, rdf:type, Y).
- type^{fk} (X, Y) :- triple^{fff} (X, P, Z), triple^{fbk} (P, rdfs:domain, Y).
- type^{fk} (X, Y) :- triple^{fff} (Z, P, X), triple^{fbk} (P, rdfs:range, Y).
- type^{fk} (X, Y) :- type^{ff} (X, Z), triple^{fbk} (Z, rdfs:subClassOf, Y).



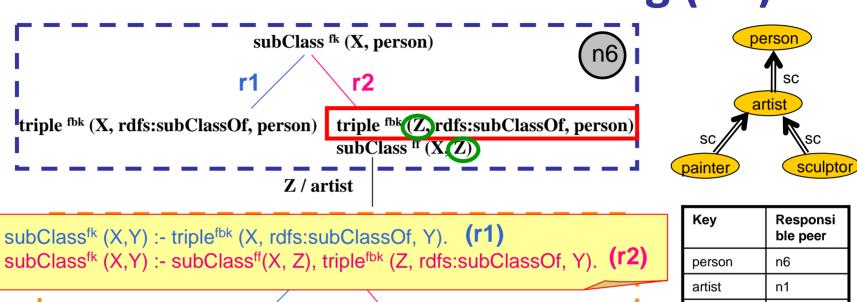


Key	Responsi ble peer
person	n6
artist	n1
painter	n3
sculptor	n4



```
 subClass^{fk} (X,Y) := triple^{fbk} (X, rdfs:subClassOf, Y). \begin{tabular}{l} (\textbf{r1}) \\ subClass^{fk} (X,Y) := subClass^{ff}(X,Z), triple^{fbk} (Z, rdfs:subClassOf, Y). \begin{tabular}{l} (\textbf{r2}) \\ \hline \end{tabular}
```

Key	Responsi ble peer
person	n6
artist	n1
painter	n3
sculptor	n4



triple fbk (X, rdfs:subClassOf, artist) subClass ff (X, Z)

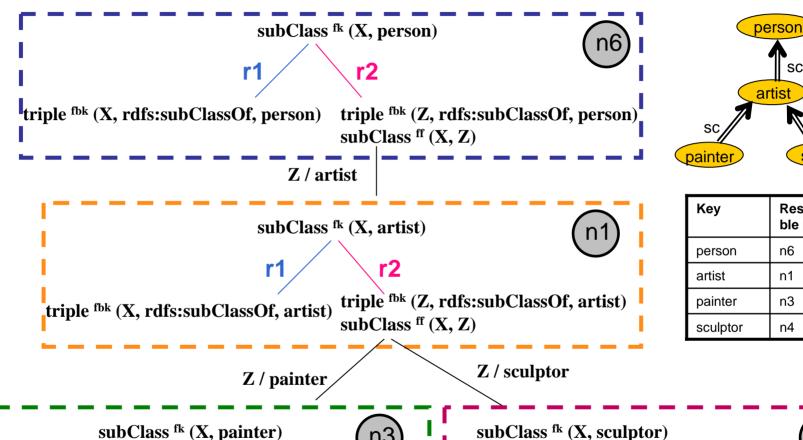
Z / painter

Z / sculptor

triple fbk (Z, rdfs:subClassOf, artist)

Key	Responsi ble peer
person	n6
artist	n1
painter	n3
sculptor	n4

Distributed Backward Chaining (BC)



Key	Responsi ble peer
person	n6
artist	n1
painter	n3
sculptor	n4

sculptor

subClass fk (X, painter) n3 **r2** riple ^{fbk} (X, rdfs:subĆlassOf, painter) triple ^{fbk} (Z, rdfs:subClassOf, painter) subClass ff (X, Z)

triple fbk (X, rdfs:subClassOf, sculptor) triple fbk (Z, rdfs:subClassOf, sculptor

subClass ff (X, Z)

- Introduction
- The system Atlas
- RDFS Reasoning in P2P networks
 - Bottom-up approach
 - Top-down approach
 - Optimized bottom-up approach
 - Evaluation
- SPARQL Query Processing and Optimization
- Conclusions

Forward chaining for magic rules (MS)

- Magic sets transformation technique
 - Rewrite rules given a specific query -
 - Bottom-up evaluation (forward chaining)
 - No unnecessary information is inferred

	Rule	Head	Body
	1	n_newTriple(P)	sup ₀₁ (P, Y)
	2	m_newTriple(P)	sup ₇₁ (P, Y)
	3	m_subProperty(Z)	$\sup_{z_1}(Z,Y)$
	4	m_subProperty(P)	n_newTriple(P)
	5	n_subClass(Z)	sup ₄₁ (Z,Y)
	6	n_type(Z)	sup ₅₁ (Z,Y)
	7	sup ₂₁ (Z, Y)	<pre>n_subProperty(Y), triple(Z, sp, Y)</pre>
	8	sup ₃₁ (P1, P)	<pre>n_newTriple(P), subProperty(Pi, P)</pre>
-	9	sup ₄₁ (Z, Y)	<pre>n_subClass (Y), triple(Z, sc, Y)</pre>
	10	sup ₅₁ (Z, Y)	<pre>m_type(Y), triple(Z, sc, Y)</pre>
	11	sup ₆₁ (P, Y)	<pre>m_type(Y), triple(P, dom, Y)</pre>
	12	sup ₇₁ (P, Υ)	<pre>n_type(Y), triple(P, range, Y)</pre>
	13	newTriple(X, P ,Y)	<pre>n_newTriple(P), triple(X, P, Y)</pre>
	14	subProperty(X, Y)	<pre>n_subProperty(Y), triple(X, sp, Y)</pre>
	15	<pre>subProperty(X, Y)</pre>	$\sup_{Z} (Z,Y)$, $\sup_{Z} Property(X,Z)$
	16	newTriple(X, P, Y)	<pre>sup₃₁(Pi,P), triple(X,Pi,Y)</pre>
	17	subClass(X, Y)	<pre>n_subClass(Y), triple(X, sc, Y)</pre>
	18	subClass(X, Y)	<pre>sup₄₁(Z,Y), subClass(X, Z)</pre>
	19	type(X, Y)	<pre>m_type(Y), triple(X, type, Y)</pre>
	20	type(X, Y)	$\sup_{\Sigma} (Z, Y), type(X, Z)$
	21	type(X, Y)	<pre>sup₀₁(P, Y), newTriple(X, P, Z)</pre>
	22	type(X, Y)	<pre>sup₇₁(P, Y), newTriple(Z, P, X)</pre>
	23	m_type(a)	

	sc	st
	painter	sculptor
sc //	sc	type
flemish	cubist	michelangelo
1	A	
type	type /	
	/	
rembrandt		
	picasso	
		Dictributed

Query: Find all painters q = (?x, rdf:type, painter)

- Introduction
- The system Atlas
- RDFS Reasoning in P2P networks
 - Bottom-up approach
 - Top-down approach
 - Optimized bottom-up approach
 - Evaluation
- SPARQL Query Processing and Optimization
- Conclusions

Experimental setup

 Algorithms have been implemented in Atlas system using Bamboo DHT [Rhea04]

Testbeds

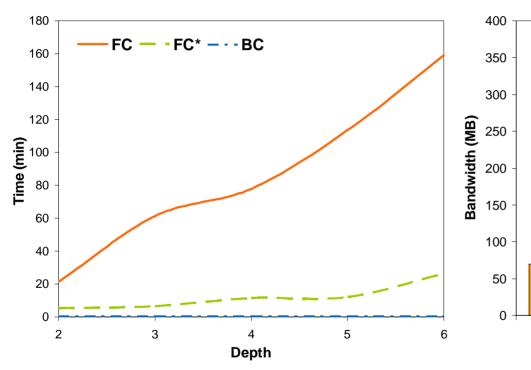
- PlanetLab (~210 nodes available at the time of the experiments)
- local cluster (30 machines 4 Atlas nodes per machine)

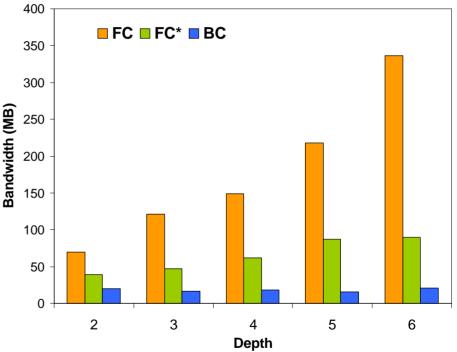
Datasets

- Synthetic data from RBench generator [Theoharis05]
 - Number of instances: 10,000 to 1M
 - RDFS class hierarchy tree depth: 2-6 (7 to 128 classes)
 - Query: Give me the instances of the root class
- LUBM benchmark
 - Number of triples: ~110,000 to ~2,700,00
 - Query: Give me the instances of a class

Store time and network traffic

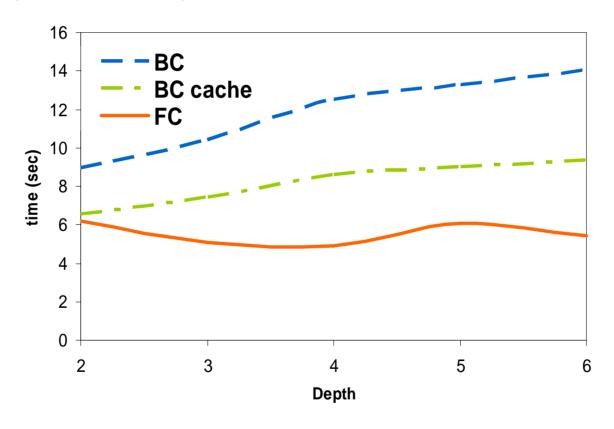
PlanetLab, RBench with 10,000 instances





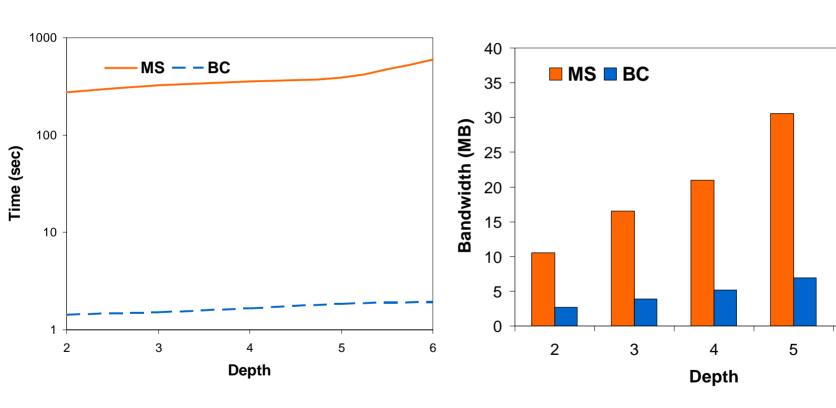
Query response time

PlanetLab, RBench with 10,000 instances



Backward chaining vs. magic sets

Cluster: RBench with 1M instances



6

- Introduction
- The system Atlas
- RDFS Reasoning in P2P networks
- SPARQL Query Processing and Optimization

Conclusions

SPARQL

SPARQL queries of basic graph patterns

```
SELECT ?x ?y ?z
WHERE {
    ?x rdf:type ub:Student . (tp1)
    ?x ub:takesCourse ?z . (tp2)
    ?x ub:advisor ?y . (tp3)
    ?y ub:teacherOf ?z . (tp4)
}
```

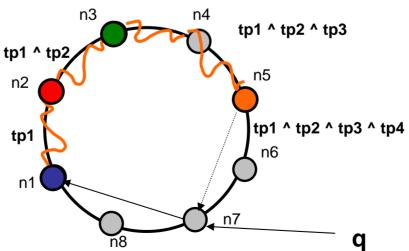
conjunctive triple patterns

```
tp1 ^ tp2 ^ tp3 ^ tp4
```

SPARQL query processing

SPARQL queries of basic graph patterns

```
SELECT ?x ?y ?z
WHERE {
    ?x rdf:type ub:Student .(tp1)
    ?x ub:takesCourse ?z . (tp2)
    ?x ub:advisor ?y . (tp3)
    ?y ub:teacherOf ?z (tp4)
}
```

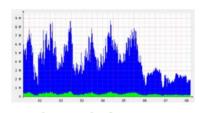


Query Optimization

 Find a query plan that optimizes the performance of a system with respect to a metric of interest



Query response time



Bandwidth consumption

Find a good ordering of the triple patterns

Minimize the size of intermediate results

Lower bandwidth consumption

Joins with smaller intermediate relations

Distributed RDF Query Processing and Reasoning in Peer-to-Peer Networks

Query Optimization

- Greedy optimization algorithms
 - selectivity-based heuristics
 - minimize the size of intermediate results

- Introduction
- The system Atlas
- RDFS Reasoning in P2P networks
- SPARQL Query Processing and Optimization
 - Selectivity estimation
 - Statistics for RDF
 - Optimization algorithms
 - Evaluation
- Conclusions

Selectivity estimation

- Single triple pattern
 - Bound-is-easier heuristic
 - Analytical estimation: $sel(tp) = sel(s) \times sel(p) \times sel(o)$

$$sel(v) = \frac{freq_c(v)}{T}$$
 frequency of value v as a component c total triples stored in the network

• Joins of triple patterns: $sel(tp_1 \wedge tp_2) = \frac{join_card(tp_1, tp_2)}{T^2}$

$$join_card(tp_1, tp_2) = \frac{T_{tp_1} \times T_{tp_2}}{\max(I_{tp_1(?x)}, I_{tp_2(?x)})} \text{ number of triples matching tp}_1, tp_2$$
 size of the domain of ?x in tp_i

- Introduction
- The system Atlas
- RDFS Reasoning in P2P networks
- SPARQL Query Processing and Optimization
 - Selectivity estimation
 - Statistics for RDF
 - Optimization algorithms
 - Evaluation
- Conclusions

Required statistics

Frequency of a triple component c with value
 v (freq_c(v))

• Size of the domain of a variable in a triple pattern $(I_{tpi(?x)})$

Frequency and variable domain size - example

(?x, advisor, ?y)

frequency of advisor \rightarrow # occurrences of advisor as a **predicate** (freq_p(advisor)) domain size of ?x \rightarrow # distinct subject values of **predicate** advisor (ds_p(advisor)) domain size of ?y \rightarrow # distinct object values of **predicate** advisor (do_p(advisor))

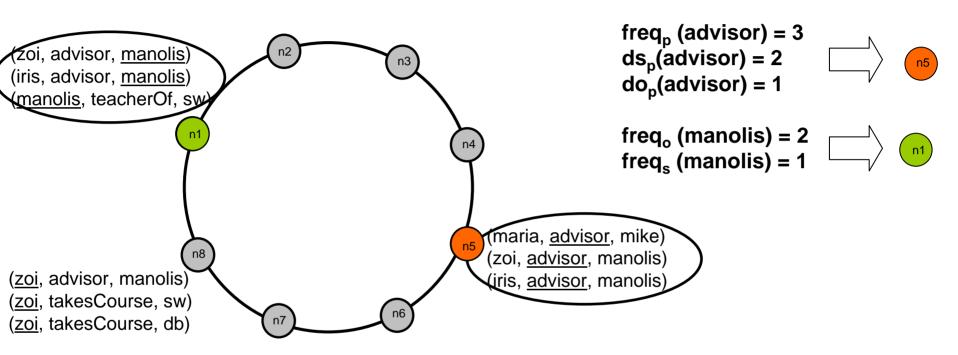
RDF triples

subject	predicate	object	
zoi	advisor	manolis	
(iris	advisor	manolis	
maria	advisor	mike	
zoi	takesCourse	db	
zoi	takesCourse	SW	
manolis	teacherOf	SW	

statistics for predicates

Value	freq _p	ds _p	do _p
advisor	3	3	2
takesCourse	2	1	2
teacherOf	1	1	1

Statistics in DHTs



Each peer is responsible for creating and maintaining statistics of its **locally** stored data and only for triple component values which are **key** at the specified peer

Statistics at each DHT peer

- Each peer creates statistics for each triple component separately
 - distinguish resource objects from class objects

subject	predicate	object	object-class
freq _s	$freq_p$	freq _o	$freq_c$
dp _s	ds_{p}	ds _o	_
dos	dp_{ρ}	dp _o	_

 Given a space budget, each peer can keep the exact frequency distribution or an estimation by using voptimal histograms

- Introduction
- The system Atlas
- RDFS Reasoning in P2P networks
- SPARQL Query Processing and Optimization
 - Selectivity estimation
 - Statistics for RDF
 - Optimization algorithms
 - Evaluation
- Conclusions

Optimization algorithms

Static optimization

- Before the query evaluation begins, at the peer that receives the query request
 - Naïve optimization algorithm (NA): Orders triple
 patterns based on their selectivity, from the most
 selective to the least selective
 - Semi-naïve optimization algorithm (SNA): Orders triple patterns based on the join selectivity between pairs of triple patterns

Optimization algorithms

Dynamic optimization

- During the query evaluation, at each peer participating in the query processing
 - Dynamic optimization algorithm (DA): join selectivity is computed at each peer participating in the query processing and estimated between the real intermediate results so far and a triple pattern
- All algorithms (static and dynamic) use a query graph (QG) representation to avoid Cartesian products

- Introduction
- The system Atlas
- RDFS Reasoning in P2P networks
- SPARQL Query Processing and Optimization
 - Selectivity estimation
 - Statistics for RDF
 - Optimization algorithms
 - Evaluation
- Conclusions

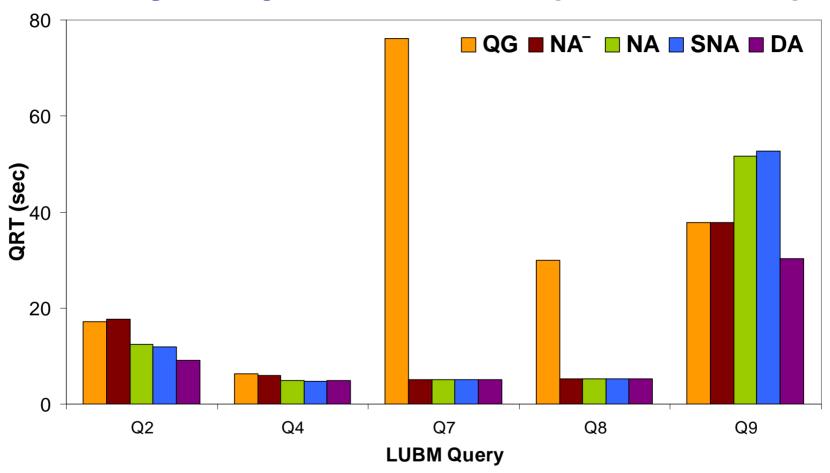
Experimental setup

- All algorithms and techniques have been implemented in Atlas
 - Mapping dictionary
- We used as a testbed PlanetLab and a local cluster server blade machines with two processors at 2.6GHz and 4GB memory 30 available machines – up to 4 nodes per machine

 up to 120 DHT peers

LUBM Benchmark (queries with more than 4 triple patterns)

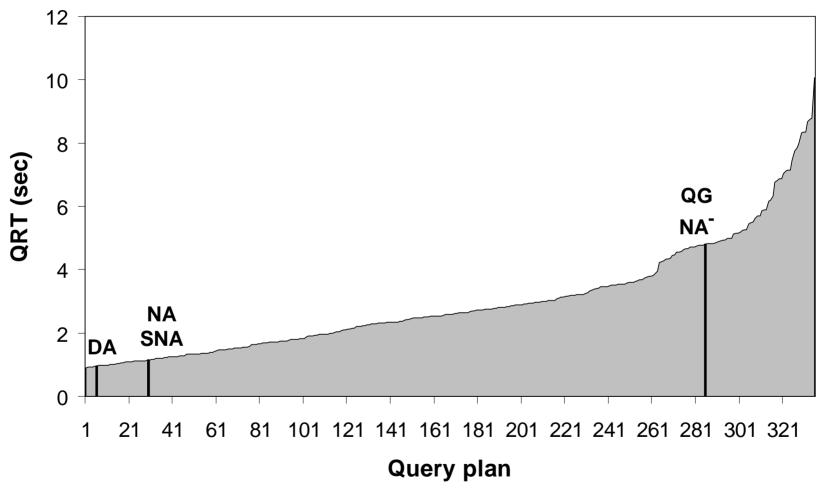
Query response time (LUBM-50)



Q7: select ?X ?Y where { ?X rdf:type ub:Student . ?Y rdf:type ub:Course . ub:AssociateProfessor147 ub:teacherOf ?Y. ?X ub:takesCourse ?Y}

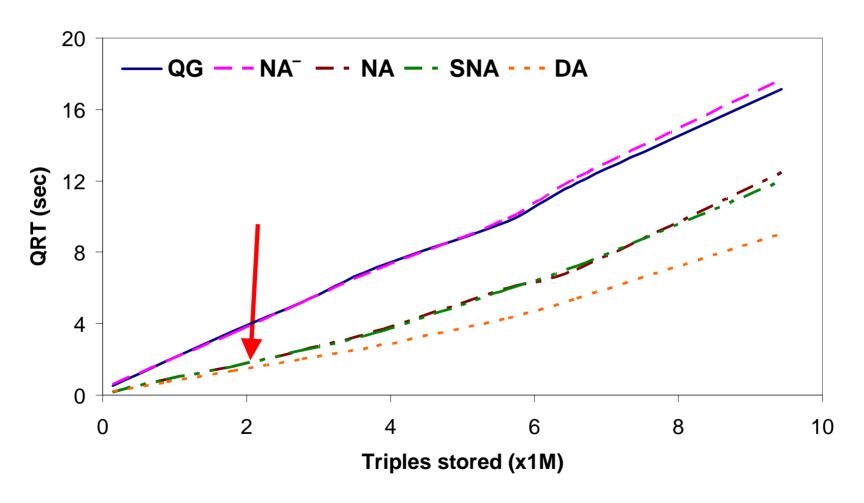
Q9: select ?X ?Y ?Z where {?X rdf:type ub:Student . ?Y rdf:type ub:Faculty . ?Z rdf:type ub:Course. ?X ub:advisor ?Y . ?X ub:takesCourse ?Z . ?Y ub:teacherOf ?Z . }

Q2 query plan space (LUBM-10)

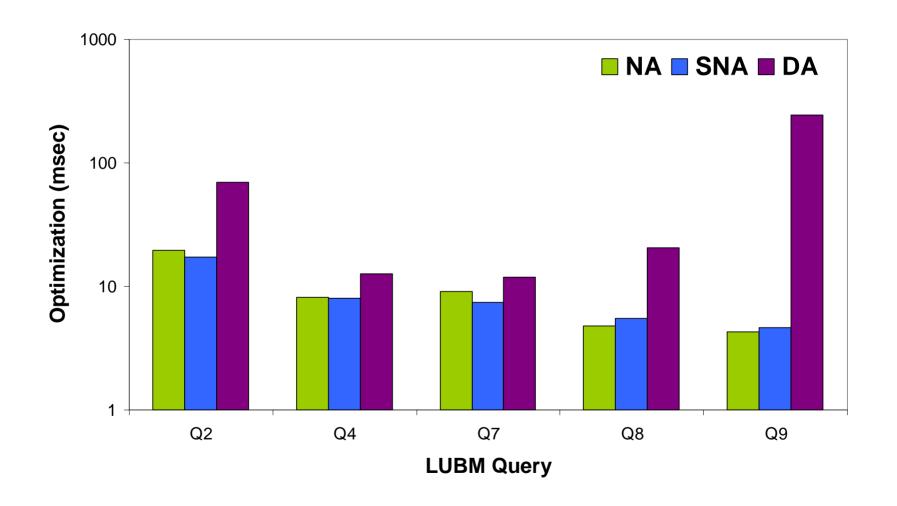


Q2: select ?X ?Y ?Z where { ?X rdf:type ub:GraduateStudent . ?X ub:undergraduateDegreeFrom ?Y . ?Y rdf:type ub:University . ?Z ub:subOrganizationOf ?Y . ?Z rdf:type ub:Department . ?X ub:memberOf ?Z . }

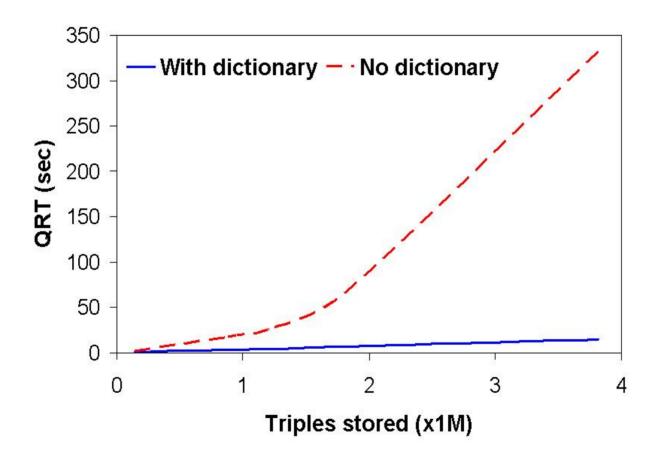
Query response time of Q2 for increasing dataset size



Optimization time (LUBM-50)



Mapping dictionary



- Introduction
- The system Atlas
- RDFS Reasoning in P2P networks

SPARQL Query Processing and Optimization

Conclusions

Conclusions

- Atlas: a P2P system for the distributed query processing and reasoning of RDF and RDFS data which is built on top of DHTs
- RDFS Reasoning on top of DHTs
 - Distributed forward chaining
 - Distributed backward chaining
 - Distributed magic sets transformation
 - Comparative study (analytically and experimentally)
 - Theoretical proofs of correctness
- SPARQL Query Processing and Optimization on top of DHTs
 - Selectivity estimation techniques
 - Statistics for RDF
 - Optimization algorithms
 - Distributed mapping dictionary

Future directions

Distributed recursive queries using Datalog

RDF in the cloud!

Thank you

Questions?

References



- **Z. Kaoudi**, K. Kyzirakos and M. Koubarakis. SPARQL Query Optimization on Top of DHTs. In *9th International Semantic Web Conference (ISWC 2010)*, Shanghai, China, November 7-11, 2010
- **Z. Kaoudi**, M. Koubarakis, K. Kyzirakos, I. Miliaraki, M. Magiridou and A. Papadakis-Pesaresi. Atlas: Storing, Updating and Querying RDF(S) Data on Top of DHTs. *Journal of Web Semantics: Science, Services and Agents on the World Wide Web*, 2010
- **Z. Kaoudi**, I. Miliaraki and M. Koubarakis. RDFS Reasoning and Query Answering on Top of DHTs. In *7th International Semantic Web Conference (ISWC 2008)*, Karlsruhe, Germany, October 26-30, 2008
- **Z. Kaoudi** and M. Koubarakis. Distributed RDFS Reasoning over Structured Overlay Networks. Extended version of ISWC 2008 submitted to *ACM TWEB*