















The problem of scale

- However, the Web is a bit larger than four websites...
 - Current LinkedDataCloud size (updated 2011): <u>~20 billion facts</u>
 - Google Knowledge's Graph: <u>25 billions of formalized facts</u> about billions of items. Volume is <u>doubling every year</u>
 - Expected outcome of OpenPhacts (large project with pharma companies): 20 - 50 billion triples, over 25 million mappings, over 20 data sources

























Problems and challenges: high I/O cost • Cause: data is distributed on several nodes and during reasoning the peers need to heavily exchange it

 Effect: hard drive or network speed become the performance bottleneck



Problems and challenges: programming complexity

- · Cause: in a parallel setting there are many technical issues to
 - handle
 - Fault tolerance
 - Data communication Execution control
 - Etc.

• Effect: Programmers need to write much more code in order to execute an application on a distributed architecture

	and the second se
-	And and the second seco
-	
	when the state has an a state of the second state of the
	The state of the local distance of the second
	PROPERTY AND ADDRESS OF THE PARTY OF THE PAR
	11 Second Control of C
	Change webs record on hits herefully.
-	Constant Annual Control of the Annual Contro
	Constants Constants and a constant and the constants of
•	Website - research by
•	ALL LL LYPE - Forest Construction
•	SALE CALS
	Senvict . allocate 'orderest' allocate

What we did

- 1. First we explored the possibility of performing forwardchaining reasoning. WebPIE
- 2. Then we researched for methods for doing backwardchaining reasoning. QueryPIE

What we did

- 1. First we explored the possibility of performing forwardchaining reasoning. WebPIE
- 2. Then we researched for methods for doing backwardchaining reasoning. QueryPIE

WebPIE

WebPIE is a forward reasoner that uses MapReduce to execute the reasoning rules

1st step:

compression

All code, documentation, tutorial etc. is available online.

http://cs.vu.nl/webpie/

WebPIE algorithm:

- Input: triples in N-Triples format
- 1) Compress the data with dictionary encoding

2) Launch reasoning

3) Decompress derived triples 2nd step: reasoning Output: triples in N-Triples format

WebPIE: Overview

RDFS reasoning

- Set of 13 rules
- All rules require at most one join between a "schema" triple and an "instance" triple
- OWL reasoning
 - Logic more complex => rules more difficult
 - The ter Horst fragment (pD*) provides a set of 23 new rules
 - Some rules require a join between instance triples
 - Some rules require multiple joins

WebPIE: Overview

RDFS reasoning

- Set of 13 rules
- All rules require at most one join between a "schema" triple and an "instance" triple
- OWL reasoning
 - Logic more complex => rules more difficult .
 - The ter Horst fragment (pD*) provides a set of 23 new rules
 - Some rules require a join between instance triples
 - Some rules require multiple joins









WebPIE: OWL reasoning

- By accepting only triples with a specific distance we avoid to derive information already derived
- General rule:
 - Every job accepts in input only triples derived in the previous two steps
 - During the execution of the nth job we derive only if:
 - The antecedent triples on the left side have distance $2^{n}(n-1)$ or $2^{n}(n-2)$
 - The antecedent triples on the right side have distance greater than 2⁽ⁿ 2)

WebPIE: Forward-chaining





What we did

- 1. First we explored the possibility of performing forwardchaining reasoning. <u>WebPIE</u>
- 2. <u>Then we researched for methods for doing backward-</u> <u>chaining reasoning. QueryPIE</u>

Backward-chaining QueryPIE: Backward-chaining · With WebPIE, we managed to perform reasoning over very Reasoning is triggered only by SPARQL queries large data · However, the intrinsic problems of performing forward-• To reduce the runtime, only the schema is pre-materialized chaining reasoning still remain does not change often Backward-chaining: · relatively small compared to the rest of the data · Advantage: we avoid to perform this (expensive) task widely used in the reasoning rules · Disadvantage: we still need to perform reasoning at query-time QueryPIE is made of two algorithms • pre-materialization algorithm: executed before the user can query the knowledge base · backward-chaining algorithm: applied when the user launches a query

QueryPIE: Backward-chaining

Pre-materialization algorithm

- <u>Goal</u>: calculate all "schema" triples
- <u>Problem</u>: cannot use forward-chaining reasoning because instance-triples can generate schema
- <u>Solution</u>: We use the *backward-chaining algorithm*:
 Problem: is incomplete if schema is not explicit
 Solution: iterate the execution until we reached completeness

QueryPIE: Backward-chaining

Backward-chaining algorithm

- Inspired by QSQ algorithm (Datalog)
- Main features
 - 1) distributes the data across cluster nodes
 - 2) prunes the reasoning tree using the inferred schema
 - 3) replicates the schema and performs most locally
- * Because of the pre-materialization, we can decrease the complexity significantly and improve the response time

QueryPIE: Backward-chaining

 Launched some example queries from the LUBM (10B triples) and LLD (5B triples)

•	Cost	pre-materialization:
•	COST	pre-materialization:

Duration Our approach Full material- ization iterations triples LUBM 1s 4d4h16m 4 390 LLD 16m 5d10h45m 7 10 million	Dataset	Reason	ng time	N.	N. derived
ization ization LUBM 1s 4d4h16m 4 390 LLD 16m 5d10h45m 7 10 million	Dataset	Our approach	Full material-	iterations	triples
LUBM 1s 4d4h16m 4 390 LLD 16m 5d10h45m 7 10 million			ization		
LLD 16m 5d10h45m 7 10 million	LUBM	1s	4d4h16m	4	390
	LLD	16m	5d10h45m	7	10 millions

QueryPIE: Backward-chaining

• Runtime of single pattern queries:

Ouomi	Runtime (ms)		Processed Triples		I/O access	
Query	Cold	Warm	Total	Output	# lookups	MB
1	299.45	6.18	5	5	43	8
2	4767.17	131.6	463	239	12	205
3	186.06	5.65	3	1	18	5
4	310.53	10.5	37	29	86	8
5	382.21	13.9	1480	719	18	4
6	3986.91	2640.81	1599987	1599987	2	12
7	470.68	224.27	0	0	670	23
8	23.85	6.33	4466	4466	1	42
9	7705.91	616.42	140	128	3540	105
10	2609.52	1163.83	28446	26860	14372	337
11	2719.40	1914.28	8546	4504	15128	64
12	1954.97	2054.66	1187944	1187944	1	10

Runtime	of the <u>L</u>	UBM SPAR	QL querie	<u>s:</u>	
	о.	Besp. tir	ne (ms.)	Results	1
	- v -	Cold	Warm	(#)	
	1	364.05	20.28	4	Ĩ
	2	368.19	25.12	6	
	3	2041.20	493.89	34	
	4	1283.88	82.42	719	
	5	2061.27	361.11	4	
	6	11536.95	2784.72	7790	
	7	1745.20	367.96	4	
	8	1664.57	74.03	224	
	9	5717.32	1163.61	15	
	10	12304.29	886.68	37118	



What have we learnt?

We stepped out from the technical contribution to identify why it "works"...

Three important lessons:

- 1st lesson: Treat schema triples differently
- $\,$ $\,$ 2^{nd} lesson: Data skew dominates the data distribution
- 3rd lesson: Certain problems appear only at a large scale

1st Lesson: treat the schema triples differently

- In the Semantic Web there is a big difference between "schema" triples and "instance" triples (A-Box vs. T-Box)
 - Schema triples: Few but very important for reasoningOther triples: Many and not always used during reasoning
 - Other triples. Many and not always used during reasoning
- Both in WebPIE and QueryPIE we replicated them on all the nodes and keep them in memory.
- Lesson: Identify the schema and treat it differently

2nd Lesson: Data skew dominates the data distribution

- Current web-data: high data skewness -> partitioning is crucial for performance
 - WebPIE: schema was replicated
- QueryPIE: data was range-partitioned (good for small queries, bad for large queries)
- Lesson: Don't ignore data skewness. Otherwise your application will not scale to large numbers

3rd Lesson: Certain problems appear only at a large scale

- · Web-scale reasoning requires high engineering efforts
 - WebPIE: ~15000 lines of code
 - QueryPIE: ~27000 lines of code
- * Which programming language to use, libraries or compression techniques are crucial questions to reach certain performance.
- Lesson: Simple proof-of-concepts prototypes are not representative

Forward vs. Backward

- Both WebPIE and QueryPIE can scale to more than <u>three times</u> <u>the size of the semantic web</u>
- Problem: Both approaches are still fragile => with some bad data the derivation can explode
- Solution: <u>approximation</u> and <u>cost models</u>

"Reasoning on a web-scale is possible. We only started to discover it."