## Materialized views for P2P XML warehousing

#### Ioana Manolescu<sup>1</sup> Spyros Zoupanos<sup>1</sup>

<sup>1</sup>GEMO/IASI group, INRIA Saclay – Île-de-France

#### July 2009 Dataring project meeting, Nantes



INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

centre de recherche SACLAY - ÎLE-DE-FRANCE



# Outline



# Introduction



# Algebraic query rewriting

- Tree pattern dialect and pattern equivalence
- Algebraic rewriting & operators
- 8 Rewriting algorithms
  - View management
    - View materialization
    - View definitions index/lookup
- 5 Experiments
- 6 Conclusion



# ViP2P at a glance

An XML data management system based on a DHT

- Declare tree pattern XML views over the network data
- Fill in the views with XML data
- Answer tree pattern queries using the existing views ۰
  - View definition lookup
  - Query rewriting  $\Rightarrow$  logical plan •
  - Translation to a (distributed) physical plan
  - Execution of the physical plan

















When *q* arrives:

- view definition lookup
- rewriting
- execution of physical plan







When *d* arrives:

- search view definitions for which v<sub>i</sub>(d) ≠ Ø
- compute  $v_i(d)$



When *d* arrives:

- search view definitions for which v<sub>i</sub>(d) ≠ Ø
- compute  $v_i(d)$
- send results

Conclusion

# **Tree pattern dialect** *P*

**a**cont



a<sub>id,cont</sub> a<sub>cont</sub>



acont aid, cont aid, val



**a**cont aid,cont aid,val

> a<sub>id</sub> b<sub>val</sub>











$a_{val}$	bcont	t <sub>id</sub>
some text	<b><c><d></d><e>some</e></c></b>	(7,6)





a <sub>val</sub>	bcont	f <sub>id</sub>
some text	<b><c><d></d><e>some</e></c></b>	(7,6)
some text	<b><b><g>text</g></b></b>	(7,6)





a <sub>val</sub>	bcont	f <sub>id</sub>
some text	<b><c><d></d><e>some</e></c></b>	(7,6)
some text	<b><b><g>text</g></b></b>	(7,6)
some text	<b><g>text</g></b>	(7,6)





• Document d, pattern  $p \Rightarrow p(d)$ 

• 
$$p_1 \equiv p_2 \text{ iff } \forall d,$$
  
 $p_1(d) = p_2(d)$ 

a <sub>val</sub>	bcont	f <sub>id</sub>
some text	<b><c><d></d><e>some</e></c></b>	(7,6)
some text	<b><b><g>text</g></b></b>	(7,6)
some text	<b><g>text</g></b>	(7,6)

# Algebraic rewriting & operators

Let  $q \in \mathcal{P}$  be a query and  $\mathcal{V} = \{v_1, v_2, \dots, v_k\}$  a set of views. A **rewriting** of *q* using  $\mathcal{V}$  is an algebraic expression

 $e(v_1, v_2, \ldots, v_k)$ 

such that  $e(\mathcal{D}) = q(\mathcal{D})$  for any document set  $\mathcal{D}$ 



# **Canonical rewritings**

# Canonical expression:



# We are interested in canonical rewritings only

## **Minimal canonical rewritings**

# Canonical expression:



# We are interested in minimal canonical rewritings only

# **Rewriting example**



# **Rewriting example**



Idea:

Compute covers of the query nodes with the view nodes.

#### Idea:

Compute covers of the query nodes with the view nodes.



v

Idea:

Compute covers of the query nodes with the view nodes.



a<sub>id,cont</sub>



Idea:

Compute covers of the query nodes with the view nodes.





Idea:

Compute covers of the query nodes with the view nodes.



Idea:

Compute covers of the query nodes with the view nodes.



No rewriting

Idea:

Compute covers of the query nodes with the view nodes.



Idea:

Compute covers of the query nodes with the view nodes.



No rewriting

Idea:

Compute covers of the query nodes with the view nodes.



Idea:

Compute covers of the query nodes with the view nodes.



No rewriting

# **Rewriting algorithms**

- SE (Subset Enumeration)
  - For each new sebset, check if a rewriting can be found
  - Test minimality at the end
- ISE (Increasing Subset Enumeration)
  - Like SE but enumerates sets from the smallest to the largest
  - Finds minimal rewritings first
- SE and ISE
  - Repeat work
  - Try all possible subsets



# **Rewriting algorithms**

Bottom-up algorithms: use smaller partial rewritings to build bigger ones

DPR (Dynamic Programming Rewriting)

• Dynamic programming style

DFR (Depth First Rewriting)

• Greedy based on the biggest query coverage





Experiments Conclusion

# **Bottom-up rewriting**



Materialized views for P2P XML warehousing

18/31

Experiments Conclusion

## **Bottom-up rewriting**









# **Rewriting algorithms trade-offs**

SE, ISE, DPR and DFR are correct and complete. They produce all minimal canonical rewritings of q given V.

- Which rewritings are "good"?
  - The one which leads to the best physical plan
  - We learn this too late!
- Heuristic: a good rewriting uses the **smallest number** of views.
  - DFR typically finds fast a solution which is reasonably good
  - ISE, DPR will need more time but return better quality results. They produce rewritings towards the end of the search

## **View materialization**

- Peer p has a view v, peer p<sub>d</sub> publishes a document d
- p indexes v on the DHT by the labels of the view
- *p<sub>d</sub>* looks up the labels and keywords of *d*
  - View set S<sub>a</sub>
  - All views v such that  $v(d) \neq \emptyset$  are is  $S_a$
- $p_d$  evaluates v(d) for each  $v \in S_a$

# View indexing and lookup for query rewriting

Query *q* asked at peer  $p \Rightarrow p$  needs to find useful views

# 4 different strategies

- Label indexing (LI):
  - Index v by each v node label
  - Look up by all node labels of q
- Return label indexing (RLI):
  - Index v by the labels of all v nodes which project some attributes
  - Same as for LI: use the labels of all q nodes as lookup keys

# Indexing and lookup view definitions

## • Leaf path indexing (LPI):

- Let LP(v) be the set of all the distinct root-to-leaf label paths of v. Index v using each element of LP(v) as key
- Look up details in the paper

#### • Return Path Indexing (RPI):

- Let RP(v) be the set of all rooted paths in v which end in a node that returns some attribute. Index v using each element of LP(v) as key
- Look up is the same as for LPI

# ViP2P platform

- Fully implemented using Java 6
- Used Berkeley DB (version 3.3.75) to store view data
- Used FreePastry (version 2.1) as our DHT network
- Experiments carried on Grid5000 using 250 machines
- 1000 ViP2P peers were deployed

## View look up performance

# We used 1440 views related to but different from query q



# View look up performance

## We used 1440 views related to but different from query q



# View building

For the experiments we indexed 2000 XMark documents and 500 views (70 related to the documents)



#### Performance of rewriting algorithms



Materialized views for P2P XML warehousing

27/31

## **Query execution**



# **Related work**

Indexing documents in the DHT [GWJD03, BC06, SHA05, AMP<sup>+</sup>08]

XPath query rewriting [BOB+04, XO05, CDO08, TYÖ+08]

- XPath: wildcard \*, union
- Rewritings: intersection, navigations, joins

Rewriting with structural constrains [ABMP07]

- Centralized setting
- Dataguide [GW97] constraints

# Summing up

- ViP2P: data access support structures for DHT based XML data management
- All the presented algorithms have been fully implemented in a functional Java based platform
- Presented at DataX 2009 (no proceedings)
- Extended version submitted for publication
- Visit us at vip2p.saclay.inria.fr!

Introduction Algebraic query rewriting Rewriting algorithms View management Experiments Conclusion

# Thank you!

Introduction Algebraic query rewriting Rewriting algorithms View management Experiments

[ABMP07] Andrei Arion, Véronique Benzaken, Ioana Manolescu, and Yannis Papakonstantinou. Structured materialized views for XML queries. In *VLDB*, pages 87–98, 2007.

[AMP+08] S. Abiteboul, I. Manolescu, N. Polyzotis, N. Preda, and C. Sun. XML processing in DHT networks. In *ICDE*, 2008.

[BC06] Angela Bonifati and Alfredo Cuzzocrea. Storing and retrieving XPath fragments in structured P2P networks. Data Knowl. Eng., 59(2), 2006.

[BOB<sup>+</sup>04] **A. Balmin, F. Ozcan, K. Beyer, R. Cochrane, and H. Pirahesh.** 

A framework for using materialized XPath views in XML query processing. In *VLDB*, 2004.

Materialized views for P2P XML warehousing

Conclusion

Introduction Algebraic query rewriting Rewriting algorithms View management Experiments

[CDO08] Bogdan Cautis, Alin Deutsch, and Nicola Onose.

Xpath rewriting using multiple views: Achieving completeness and efficiency. In *WebDB*, 2008.

[GW97] **Roy Goldman and Jennifer Widom.** Dataguides: Enabling query formulation and optimization in semistructured databases. In *VLDB*, 1997.

[GWJD03] L. Galanis, Y. Wang, S.R. Jeffery, and D.J. DeWitt. Locating data sources in large distributed systems. In *VLDB*, 2003.

[SHA05] Gleb Skobeltsyn, Manfred Hauswirth, and Karl Aberer. Efficient processing of XPath queries with structured overlay networks.

Materialized views for P2P XML warehousing

Conclusion

 Introduction
 Algebraic query rewriting
 Rewriting algorithms
 View management
 Experiments
 Conclusion

 In
 OTM Conferences (2), 2005.

 [TYÖ+08]
 Nan Tang, Jeffrey Xu Yu, M. Tamer Özsu, Byron<br/>Choi, and Kam-Fai Wong.<br/>Multiple materialized view selection for xpath query<br/>rewriting.<br/>In ICDE, pages 873–882, 2008.

 [XO05]
 W. Xu and M. Ozsoyoglu.

Rewriting XPath queries using materialized views. In *VLDB*, 2005.

