

Path Query Routing in Unstructured Peer-to-Peer Networks

Nicolas Bonnel, Gildas Ménier, Pierre-francois Marteau
Laboratoire Valoria - Université de Bretagne Sud

August 29, 2007



VALORIA





- 1 Introduction
 - Context
 - Overview
- 2 Related work
 - P2P architecture
 - Bloom filters
- 3 Content based routing of path queries
 - Multi Level EDBF
 - Clustering
 - Preliminary results
- 4 Conclusion

Context

Context

- Indexing very large database
- Semi-structured information (XML)
- Need to index the structure of documents
- Need to answer approximative queries



Queries

- Exact : article/title
- 1 unknown element : article/?/paragraph
- 0 or more unknown elements : */paragraph

Architecture

Overview

- Distributed XML database
- The system constrains the location and replication of data

Resources scavenging


- Allow to use more computers, cheap cost
- Ex : SETI

Peer to Peer architecture

- Fault tolerance
- Scalability

- 1 Introduction
 - Context
 - Overview
- 2 **Related work**
 - P2P architecture
 - Bloom filters
- 3 Content based routing of path queries
 - Multi Level EDBF
 - Clustering
 - Preliminary results
- 4 Conclusion

Structured p2p network



Chord, CAN, Tapestry, ...


Advantage

- Easy to retrieve rare items

Limitations

- Approximatives and ranged queries very costly
- Load balancing problems

Unstructured p2p network



Gnutella [Clip2, 2002], ...

Advantages

- Highly replicated items can be retrieved at a cheap cost
- Can control data placement


Limitation

- Very costly to retrieve rare items

Bloom Filters [Bloom, 70]

Definition

- A array of m bits.
- $h_i : 0 \leq i < k$ k hash functions.
- $\text{insert}(x) : \forall i : A[h_i(x)] = 1$
- $\text{query}(x) : \text{true if}$
 $\forall i : A[h_i(x)] = 1$



False positives

- False positives are possible, but false negatives are not
- Probability of false positive : $(1 - (1 - \frac{1}{m})^{kn})^k$


Exponentially Decaying Bloom Filter [Kumar, 2005]

EDBF

- $\theta(x) = |\forall i : A[h_i(x)] == 1|$
- Can be used to encode stochastic routing tables.
- $\theta(x)/k$: probability to find x among a specific link.
- n hop count from element : $\theta(x)/k = 1/d^n$


Update

- Copy filter of each neighbor
- Bits of the copy are set to 0 with a probability $(1 - 1/d)$
- OR with local filter
- ex : propagation with a decay $d = 2$



Multi Level Bloom Filter [Koloniari, 2004]

article/body/chapter




Breadth Bloom Filter

- Set of Bloom filters
- Each element at level i is stored in BBF_i

Depth Bloom Filter

- Set of Bloom filters
- Each XML path of length i is stored in DBF_i

article/body/chapter



- 1 Introduction
 - Context
 - Overview
- 2 Related work
 - P2P architecture
 - Bloom filters
- 3 Content based routing of path queries
 - Multi Level EDBF
 - Clustering
 - Preliminary results
- 4 Conclusion

Multi Level EDBF

Breadth EDBF

- Exponential decaying version of BBF
- Additional filter to store elements in a reverse order

article/body/chapter



chapter



body



article


Querying

- Product of probabilities
- BBF can answer $E/*$ queries
- RBBF can answer $*/E$ queries
- Both filters can answer $E/* / E$ queries

Data clustering

Agent

- An agent carries an indexed XML path chosen at random
- Moves randomly on the network
- If better node \rightarrow moves the indexed XML path on this node
- Comparison function : number of XML path's elements in the filter of the node and its neighborhood



Example

- BBF_1 contains A
- BBF_2 contains B
- $RBBF_1$ contains C
- $RBBF_2$ contains B
- path $A/B/C$ have a score of 4

Experiments settings

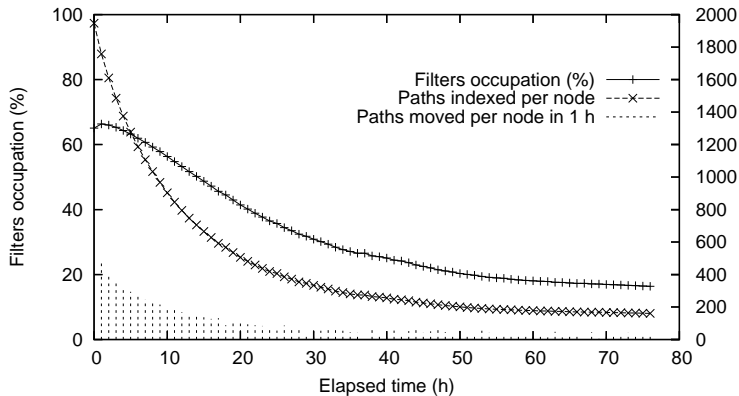
Network topology

- 200 nodes
- Random graph
- Node degree between 3 and 8


Settings

- 260 000 XML documents from Wikipedia (1.5 GByte)
- Filter's size : 8192 (2^{13}), 3 filters in sets
($BBF_1, BBF_2, BBF_3, RBBF_1, RBBF_2, RBBF_3$)
- Number of hash functions : 32
- 1000 queries generated at random

Preliminary results



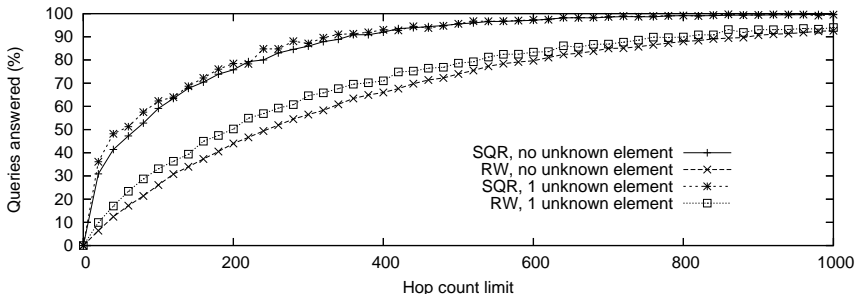
Preliminary results



2 element queries

- No unknown element : article/title, section/paragraph, ...
- 1 unknown element : ?/abstract, article/?, ...


Preliminary results



3 element queries

- No unknown element : article/section/paragraph, ...
- 1 unknown element :
article/?/paragraph, ?/section/paragraph, ...

Preliminary results



Ancestor-descendant queries

- article/*/paragraph, article/*/abstract

- 1 Introduction
 - Context
 - Overview
- 2 Related work
 - P2P architecture
 - Bloom filters
- 3 Content based routing of path queries
 - Multi Level EDBF
 - Clustering
 - Preliminary results
- 4 Conclusion

Conclusion

Contribution

- Routing of approximative XML path queries
- Data clustering of path indexes.

Experiments

- Increase routing performances compared to random walk
- Good performances with rare elements





Future Work

- Larger network
- Information replication
- Take into account element attributes

Acknowledgements

This research was supported by Region Bretagne.

References

-  [Burton H. Bloom.](#)
Space/time trade-offs in hash coding with allowable errors.
Communications of the ACM, 13(7) :422–426, 1970.
-  [Clip2.](#)
The gnutella protocol specification v0.4, 2002.
-  [G. Koloniari and E. Pitoura.](#)
Content-based routing of path queries in peer-to-peer systems.
In *Proceedings of the EDBT'04 International Conference, Heraklion, Crete, Greece*, 2004.
-  [Abhishek Kumar, Jun Xu, and Ellen W. Zegura.](#)
Efficient and scalable query routing for unstructured peer-to-peer networks.
In *Proc. of IEEE Infocom*, 2005.