

Dynamic Authenticated Index Structures for Outsourced Databases

Feifei Li, Marios Hadjieleftheriou, George Kollios, Leonid Reyzin
Boston University
AT&T Labs-Research

Presenter : Nima Najafian

1

Outline

- ☺ The Model
- ☺ Motivation
- ☹ Problem
- ☺ Solution
- ☺ Background
- ☺ Papers contributions
- ☺ Experimental validation

Outsourced Database Model

Owner: publish data
Servers: host the data and provide query services
Clients: query the owner's data through servers

clients servers owner

3

Motivation

- Advantages
 - The data owner does not need the hardware / software / personnel to run a DBMS
 - The owner achieves economies of scale
 - The client enjoys better quality of service
- A main challenge
 - The service provider is not trusted, and may return incorrect query results

Problem

☹ Un-trusted Servers

Un-trusted server

- Lazy: incentives to perform less
- Curious: incentives to acquire information
- Malicious 🤪
 - Incorrect results (could be bugs)
 - Possibly compromised

Problem 1: Injection

Select * from T where 5<A<11

client Returns 7, 8, 9

owner

	A	B
r ₁	...	
...	...	
r ₁₋₁	4	
r ₁	7	
r ₁₋₁	9	
r ₁₋₂	11	

server

7

Problem 2: Drop

Select * from T where 5<A<11

client Returns 7

owner

	A	B
r ₁	...	
...	...	
r ₁₋₁	4	
r ₁	7	
r ₁₋₁	9	
r ₁₋₂	11	

server

8

Solution

- ☺ The Model
- ☺ Motivation
- ☹ Problem
- ☺ Ability to authenticate without trusting the server (Query Authentication)

Query Authentication: (the dimensions)

- Query Correctness
results do exist in the owner's database ~ injection
- Query Completeness
no records have been omitted from the result ~ drop
- Query Freshness ★
results are based on the most current version of the database (this will bring a third problem into the picture) ~omission

10

General Approach

Verification Object (VO)

Authenticated Structures

Query results

clients servers owner

11

Background


Cryptographic essentials

1: Collision-resistant hash functions

- It is computational hard to find x_1 and x_2 s.t. $h(x_1)=h(x_2)$
- Computational hard? Based on well established assumptions such as discrete logarithms
- SHA1
- Observations:
 - variable input size \rightarrow 20 bytes
 - Computation cost: 2-3 μ s (for up to 500 bytes input)
 - Storage cost: 20 bytes
 - Under Crypto++ [crypto] and OpenSSL [openssl]

13

2: Public key digital signature schemes



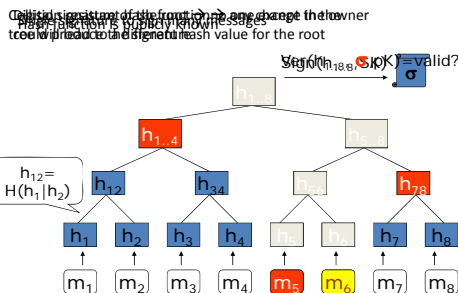
14

2: Public Key Digital Signature Schemes

- Formally defined by [GMR88]
 - The message has not been changed in any way
 - The message is indeed from the sender (corresponding to the public key)
 - No one except the secret key owner could produce a signature
- One such scheme: RSA [RSA78]
- Observations
 - Computation cost: about 3-4 ms for signing and more than 100 μ s for verifying
 - Storage cost: 128 bytes
- 3: Signature Aggregation (Condensed RSA)
 - Checking one aggregated signature is almost as fast as an individual signature

15

4: Merkle Hash Tree[M89]-Amortizing Signature Cost



16

Correctness and Completeness

- Correctness, Completeness:
 - Any change in the tree will lead to different hash
 - Relative position of values is authenticated
- Authentication:
 - Signing the root with SK

17

Contributions

☺ Proposed authenticated structures

- 🕒 Getting to know B+ trees
- 🕒 The idea of changing
- 🕒 ASB Tree (based on existing work)
- 🕒 MB tree (based on existing work)
- 🕒 EMB tree
- 🕒 Freshness (third dimension of query Authentication)

B+ - Tree Structure


- A typical node contains up to $n - 1$ search key values K_1, K_2, \dots, K_{n-1} , and n pointers P_1, P_2, \dots, P_n . The search key values are kept in sorted order.
- The pointer P_i can point to either a file record or a bucket of pointers which each point to a file record.



19


B+ - Tree File Organization

In a B+ - Tree file organization, the leaf nodes of the tree stores the actual record rather than storing pointers to records.




20

Range Authentication – A Simple Approach



21

Signature-Based Approach: ASB Tree based on [PJR05]




1. order database tuples w.r.t query attribute
2. sign consecutive pairs
3. build B+ tree on top of it
4. return tuples $[a-1, b+1]$ together with signatures in $[a-1, b]$. (query is $[a, b]$) (a, b here are index)
5. verify any two consecutive pairs

22

Condensed RSA (NDSS'04)

- Server:
 - Selects records matching posed query
 - Multiplies corresponding RSA signatures
 - Returns **single** signature to querier



23

Comparing Cryptographic OP

- one hashing takes 2-3 μ s
 - Modular Multiplication -100 times slower
 - Verifying -1000 times slower
 - Signing -10000 times slower

$$t_{\text{Hashing}} < t_{\text{mod_M}} < t_{\text{ver}} < t_{\text{Sign}}$$

24

Reduce S/C communication Cost

- Aggregation Signature: Condensed RSA

$\sigma = \text{combine}(\sigma_1, \dots, \sigma_k)$

Overhead: computation cost of modular multiplication with big modular base number, close to 100 μs

25

Signature Chaining Issues

- A heavy burden on the owner to produce the signatures
- Overhead on the client to verify the aggregated signature
- Storage overhead at the server to store the signatures (which potentially leads to higher computational cost to retrieve them)
- High communication overhead on both the server and the owner, in order to exchange the signatures

Merkle B(MB) Tree: Natural Extension for Range Query

- Use a B⁺-tree instead of a binary search tree:

- Extend it with hash information:

leaf node

27

Merkle B(MB) Tree: Natural Extension for Range Query

$h_i = H(h_{10} | \dots | h_{1r})$

For root node, $\sigma = \text{Sign}(h_0 | \dots | h_r)$

28

Extends to Range Query: $f=2$ (f is the fanout)

Select * from T where $5 < A < 11$

VO: 5, 12, $h_{1,4}$, σ

LB(q) q RB(q)

29

Client Side Verification

Select * from T where $5 < A < 11$

VO: 5, 12, $h_{1,4}$, σ

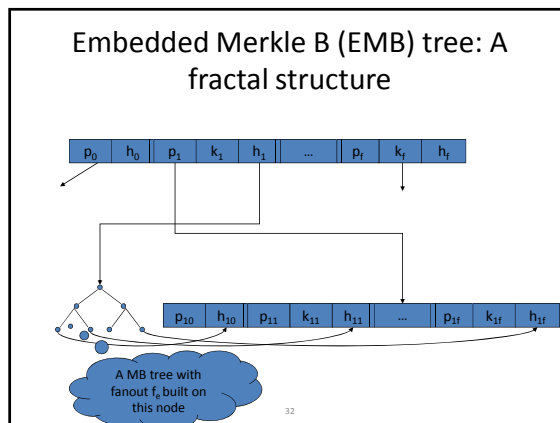
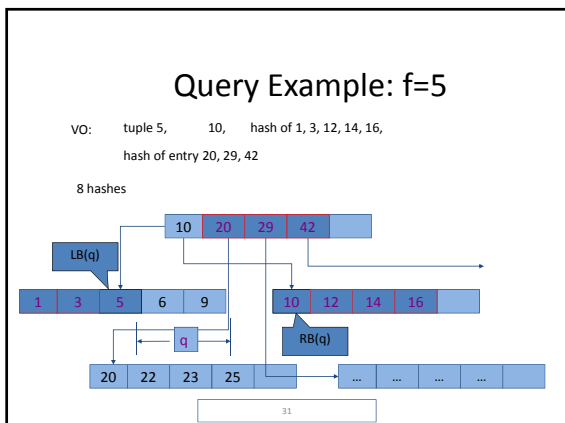
Query results: 6, 9

Valid?

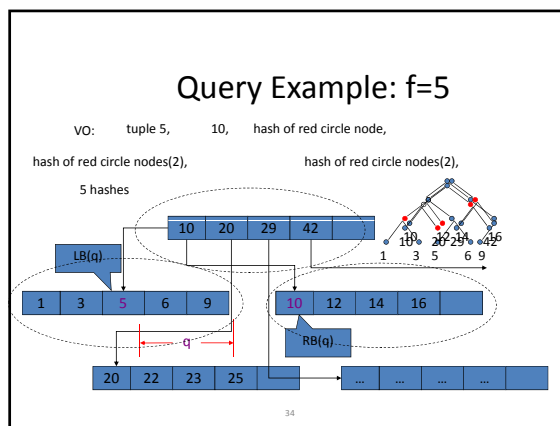
Unknown to the client

Reconstruct query subtree

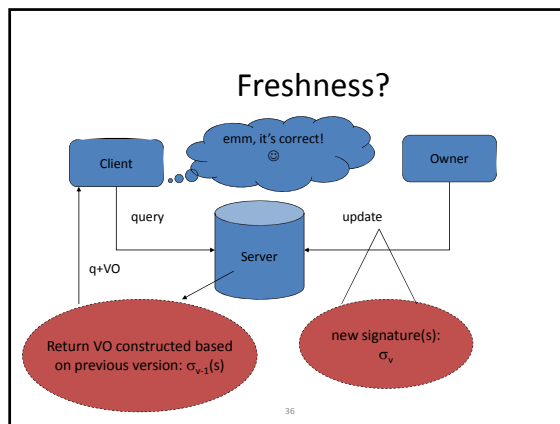
30



- ### EMB tree Analysis
- We can show that:
 - Query cost is as a MB tree with fanout f_k
 - Authentication cost (c/s comm. cost and client verification cost) is as a MB tree with fanout f_e , intuition:
 - f_k is smaller than a normal MB tree given a page size P



- ### EMB tree's variants
- Don't store the embedded tree, build it on the fly – EMB tree
 - Fanout f_k is as a normal MB tree, better query performance, better storage performance
 - Use multi-way search tree instead of B+ tree as embedded tree – EMB* tree
 - Hash path in the embedded tree could stop in index level, not necessary to go to the leaf level, hence reduce the VO size



Problem 3: Omission

Select * from T where 5 < A < 11

The diagram illustrates a client-server interaction. The client sends a query to the server. The server returns results (7, 9). The owner updates the table with new values (4, 6, 7, 9, 11).

	A	B
r_1	...	
...	...	
r_{i-1}	4	
r_i	7	
r_{i+1}	9	
r_{i+2}	11	

Update

	A	B
r_1	...	
...	...	
r_{i-1}	4	
r_i	6	
r_{i+1}	7	
r_{i+2}	9	
r_{i+3}	11	

37

Solution to Freshness

- Must have client-owner communication
 - Reduce this communication cost is the key issue
 - Observation: this cost is correlated with the number of signatures maintained in the authentication structure used by the owner

38

Other Query Types

- Join
- Projection
- Aggregate

39

Tradeoff: query vs. authentication efficiency

- Key observations:
 - Query efficiency vs. authentication efficiency
 - Impossible to have one solution that optimizes all cost metrics

40

Comparing Cryptographic OP

- one hashing takes 2-3 μ s
 - Modular Multiplication -100 times slower
 - Verifying -1000 times slower
 - Signing -10000 times slower
- Why is verifying faster?!

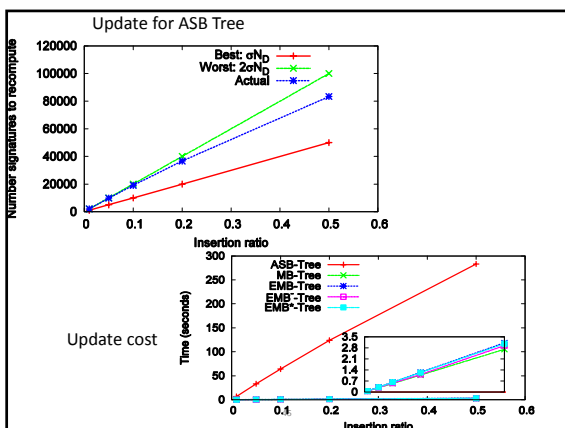
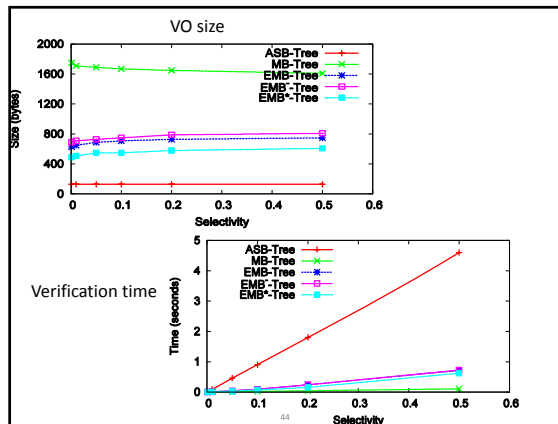
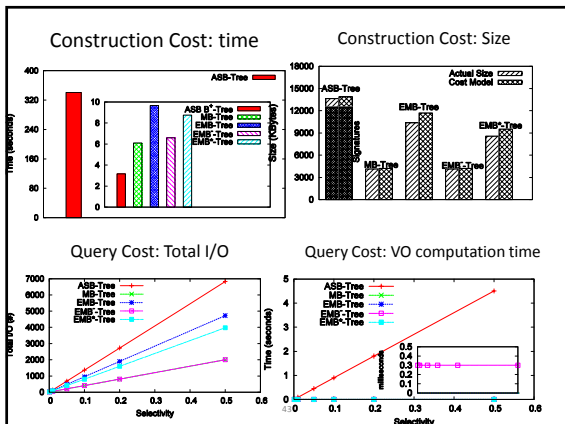
$$t_{\text{Hashing}} < t_{\text{mod_M}} < t_{\text{ver}} < t_{\text{Sign}}$$

41

Experiments

- Experiment setup
 - Crypto function – Crypto++ and OpenSSL
 - Pagesize: 1KB
 - 100,000 tuples
 - 2.8GHz Intel Pentium 4 CPU
 - Linux Machine

42



References

- [CRYPTO] Crypto++ Library. <http://www.eskimo.com/~weidai/cryptlib.html>.
- [DGMS00] P. Devanbu, M. Gertz, C. Martel, and S. G. Stubblebine. Authentic third-party data publication. In IFIP Workshop on Database Security, 2000.
- [DGMS03] P. Devanbu, M. Gertz, C. Martel, and S. Stubblebine. Authentic data publication over the internet. Journal of Computer Security, 11(3), 2003.
- [GR97] R. Gennaro, P. Rohatgi. How to Sign Digital Streams. In Crypto 97
- [GMR88] S. Goldwasser, S. Micali, and R. L. Rivest. A digital signature scheme secure against adaptive chosen-message attacks. SIAM Journal on Computing, 17(2), April 1988.
- [HIM02] H. Hacigumus, B. R. Iyer, and S. Mehrotra. Providing database as a service. In ICDE, 2002.
- [M90] K. McCurley. The discrete logarithm problem. In Cryptology and Computational Number Theory, Proc. Symposium in Applied Mathematics 42. American Mathematical Society, 1990.
- [M89] R. C. Merkle. A certified digital signature. In CRYPTO, 1989.

Thank you !