CERIAS

Ensuring Correctness over Untrusted Private Database

Sarvjeet Singh **Sunil Prabhakar**
Department of Computer Science, Purdue University
{singh35,sunil}@cs.purdue.edu

Motivation

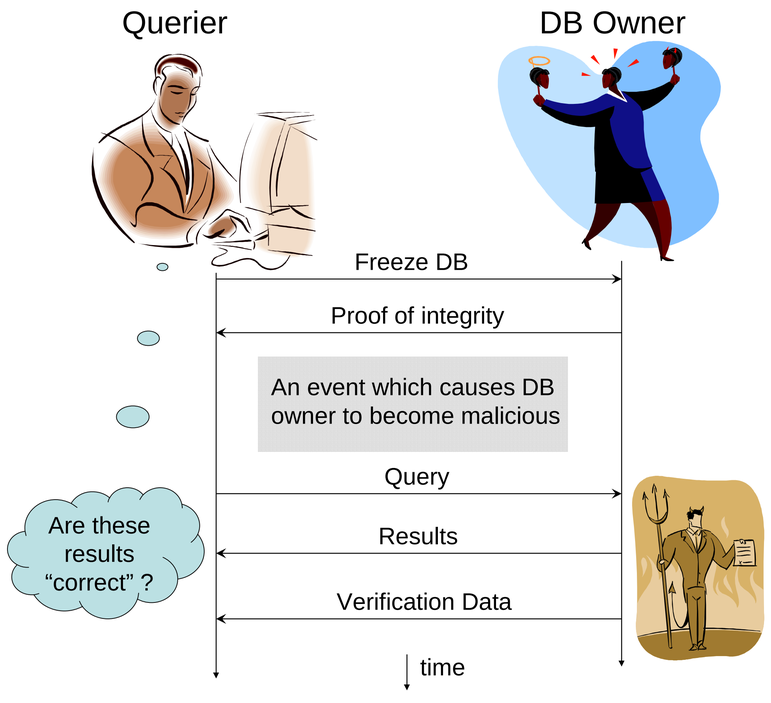

Assumptions

- Both entities do not trust each other
- The database owner is not willing to reveal anything other than the results of the query
- The Querier wants a proof that query results are not modified in response to events such as submission of query
- No restriction on how the query results can be modified

Issues

- Efficiency: Size of proof, Cost of proof generation, Size of verification object, Cost of verifying, Exposure of data
- Proof phase is very frequent compared to verification phase
- Attributes with small domain
- Granularity of hashing: Tradeoff between degree of exposure and generation cost

α - Correctness

α-correctness: Proves that the data in query result is not tampered

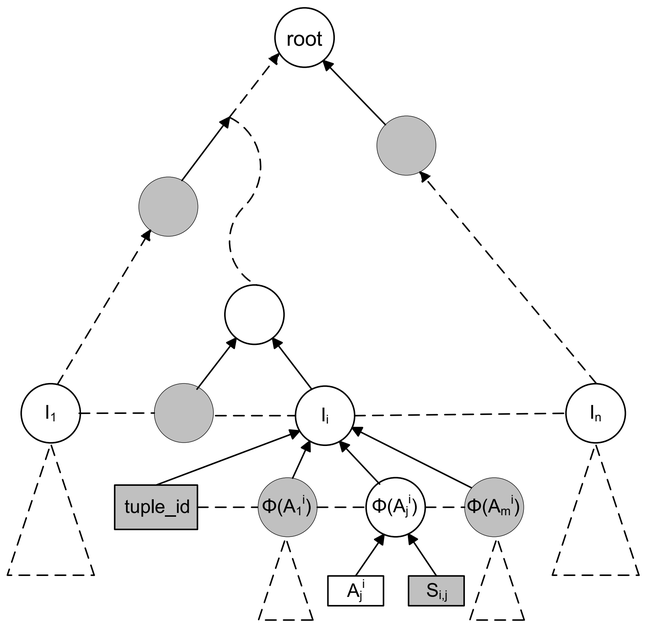

β - Correctness

β-correctness: Proves that the query operations (selections, projects, joins) are executed correctly and no tuples are missing.

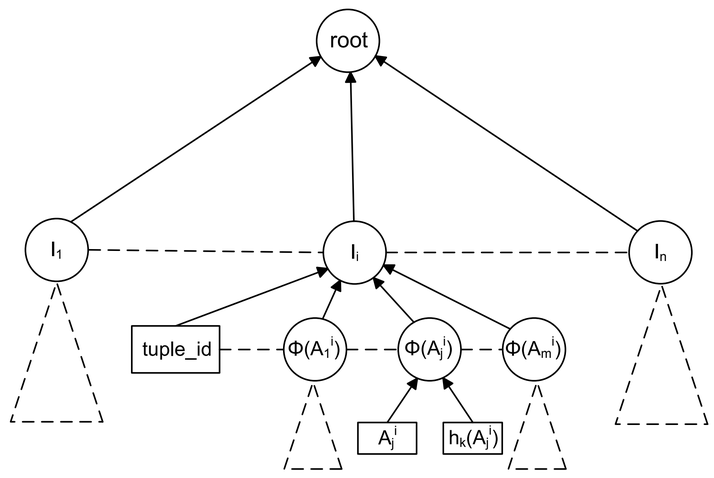

Experiments

Proposed solutions are tested through implementation using PostgreSQL and real data. The results show that they are easy to implement and overheads are acceptable

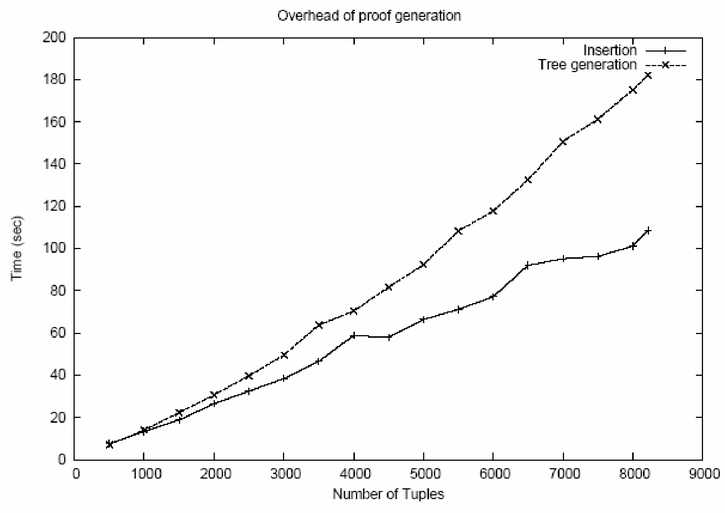

PURDUE UNIVERSITY

CERIAS

Discovery Park e-Enterprise Center