IntegriDB: Verifiable SQL for Outsourced Databases

Yupeng Zhang, Jonathan Katz and Charalampos Papamanthou

Verifiable Database

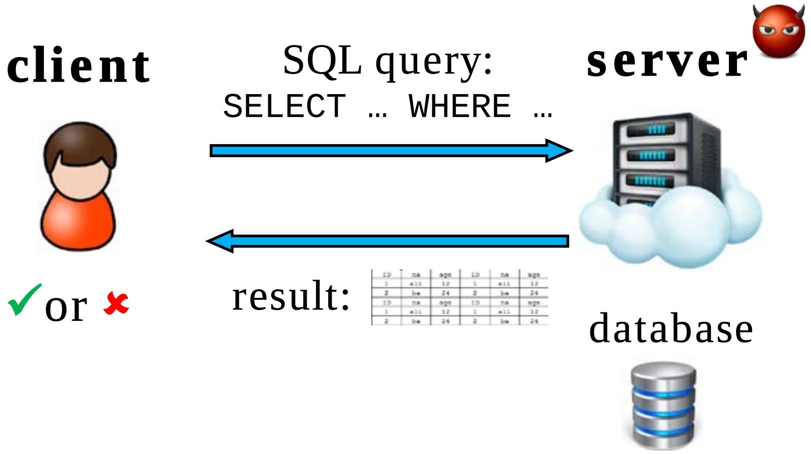

Authenticated Interval Tree

row ID	student ID	age	GPA	Name
1	747	18	3.5	Alice
2	781	24	3.3	Bob
3	715	21	3	Cathy
4	721	20	3.7	David

SELECT *
WHERE (16<*age*<23)
AND (*student ID* > 720)

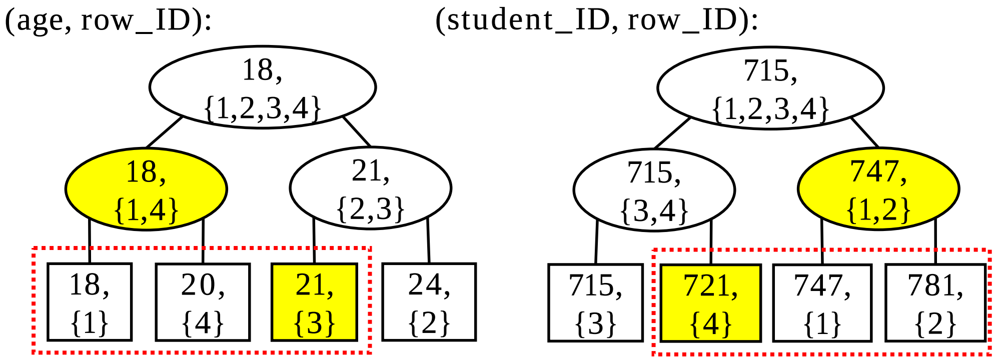

Final result: ({1,4} U {3}) ∩ ({1,2} U {4}) = {1, 4}

Set Accumulator

$$acc(A) = g^{\prod_{x\in A}(x+s)}$$

Constant size on client side. Can validate the result of set operations.

Replace sets in the authenticated interval trees with their accumulators.

Supported Queries

1. Multi-range.
2. Join: reduce to intersections.
3. Sum: modified accumulator.

$$acc(A) = g^{\prod_{x\in A}(x^{-1}+s)}$$

4. Count: reduce to sum.
5. Max/Min: reduce to range with unique answer.
6. Nested queries.
7. Updates.

Our System

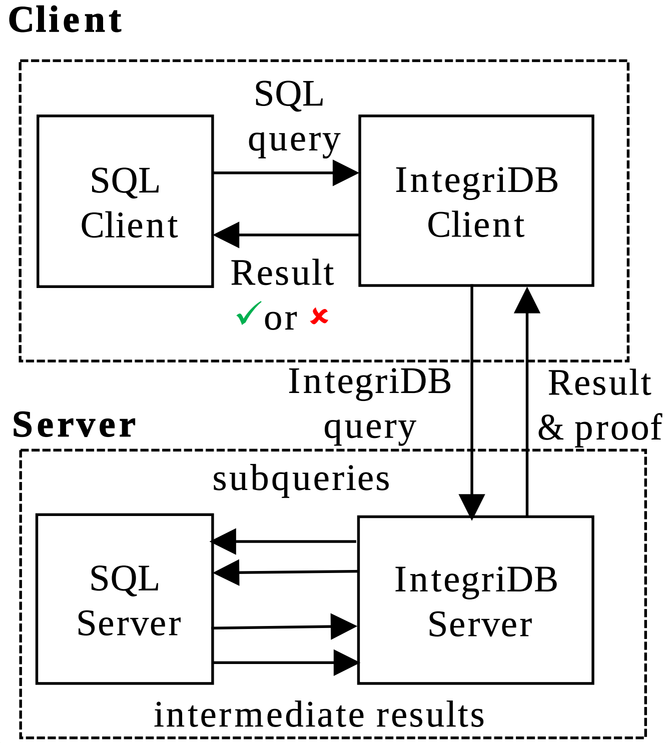

Experimental Results

TPC benchmark:

TPC-H database: largest table with 6 million rows and 16 columns (2.8GB).
TPC-H query #19: 7-dimensional range on two tables + join + sum.

Setup time	Prover time	Verification time	Proof size	Update time	Digest size
25272.76s	6422.13s	232ms	184.16KB	150ms	256bits

IntegriDB supports 12 out of 22 queries in TPC-H benchmark, 94% of the queries in TPC-C benchmark

Improvement upon prior work:

Table: 1000 rows and 10 columns. Query: 10-dimensional range + sum.

	Libsnark (circuit-based)	SNARKs for C (RAM-based)	IntegriDB
setup time	187.96s	2000s*	13.878s
prover time	47.57s	1000s*	10.420s
verification time	8ms	10ms*	112ms
proof size	288 Bytes	288 Bytes	84 KB