

**The 17th IEEE International Conference on
Trust, Security and Privacy in Computing and Communications
(IEEE TrustCom 2018)**

Query Guard: Privacy-preserving Latency-aware Query Optimization for Edge Computing

Runhua Xu, Balaji Palanisamy and James Joshi

University of Pittsburgh

Pittsburgh, PA, US

runhua.xu@pitt.edu



University of Pittsburgh
School of Computing
and Information



LERSAIS

The Laboratory for Education and Research on
Security Assured Information Systems



Edge Computing

It allows data produced by IoT devices to be processed geographically closer to where it is created instead of sending it across long routes to data centers/clouds.

(The background image is from xtelesis.com)

The Laboratory for Education and Research on
Security Assured Information Systems (LERSAIS)

Wednesday, August 1, 2018

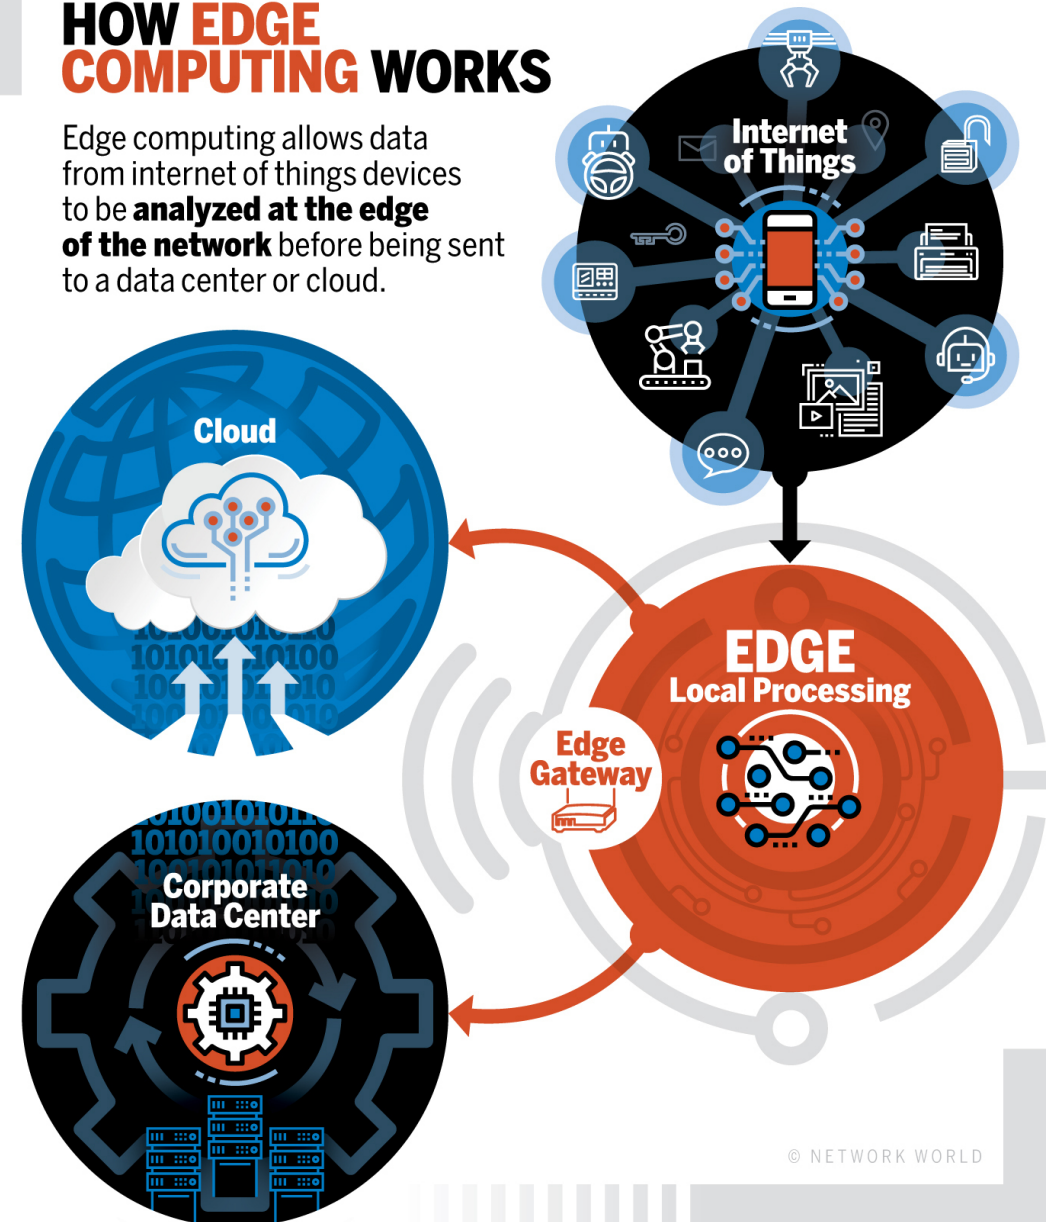
Why/How Edge Computing

- Why does edge computing matter
 - *IoT devices have poor connectivity*
 - *It's not efficient for IoT devices to be constantly connected to a central cloud.*
 - *latency-sensitive processing requirement*
- How edge computing works
 - *Triage the data locally*

Wednesday, August 1, 2018

HOW EDGE COMPUTING WORKS

Edge computing allows data from internet of things devices to be **analyzed at the edge of the network** before being sent to a data center or cloud.

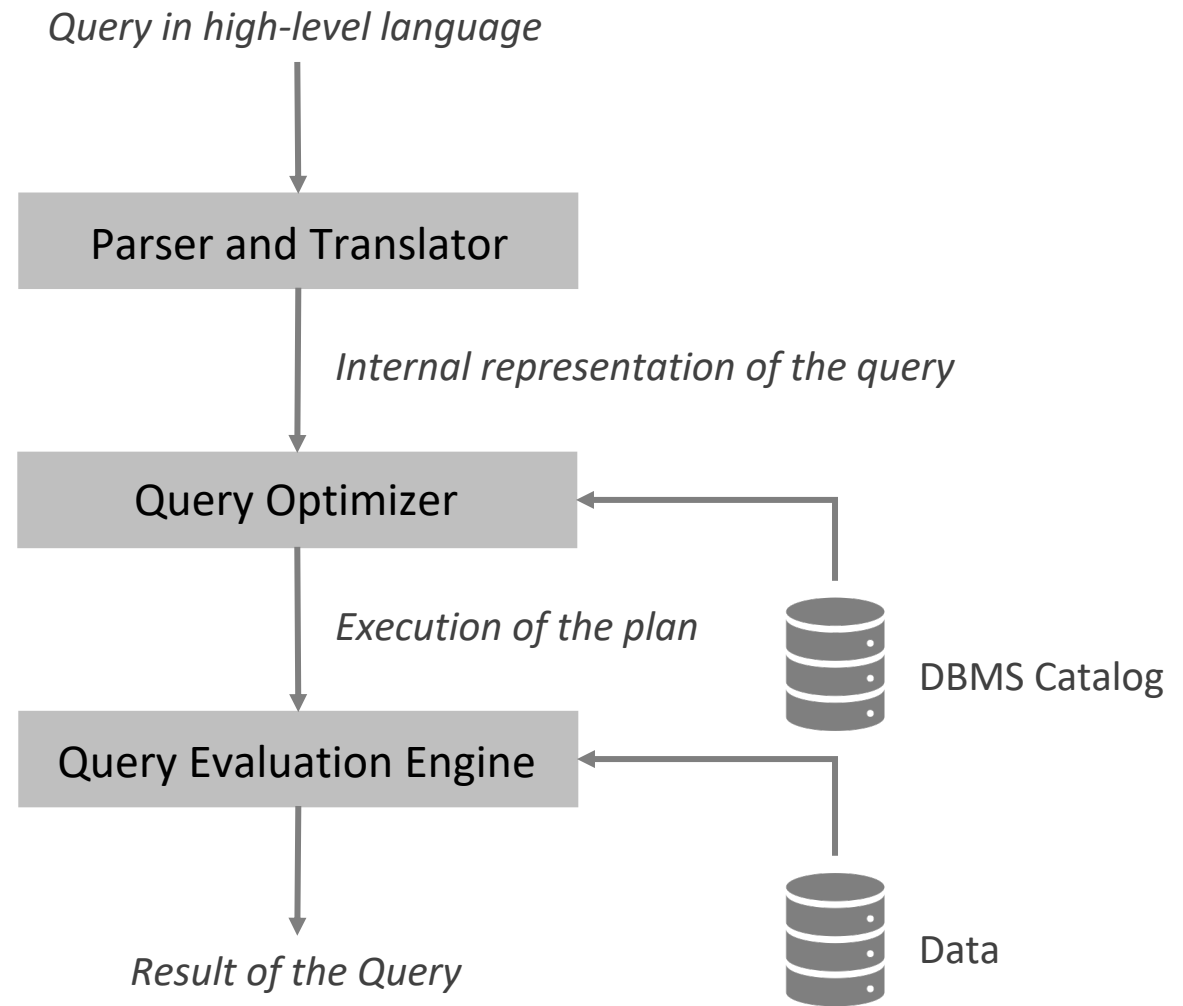


Overview of Query Processing

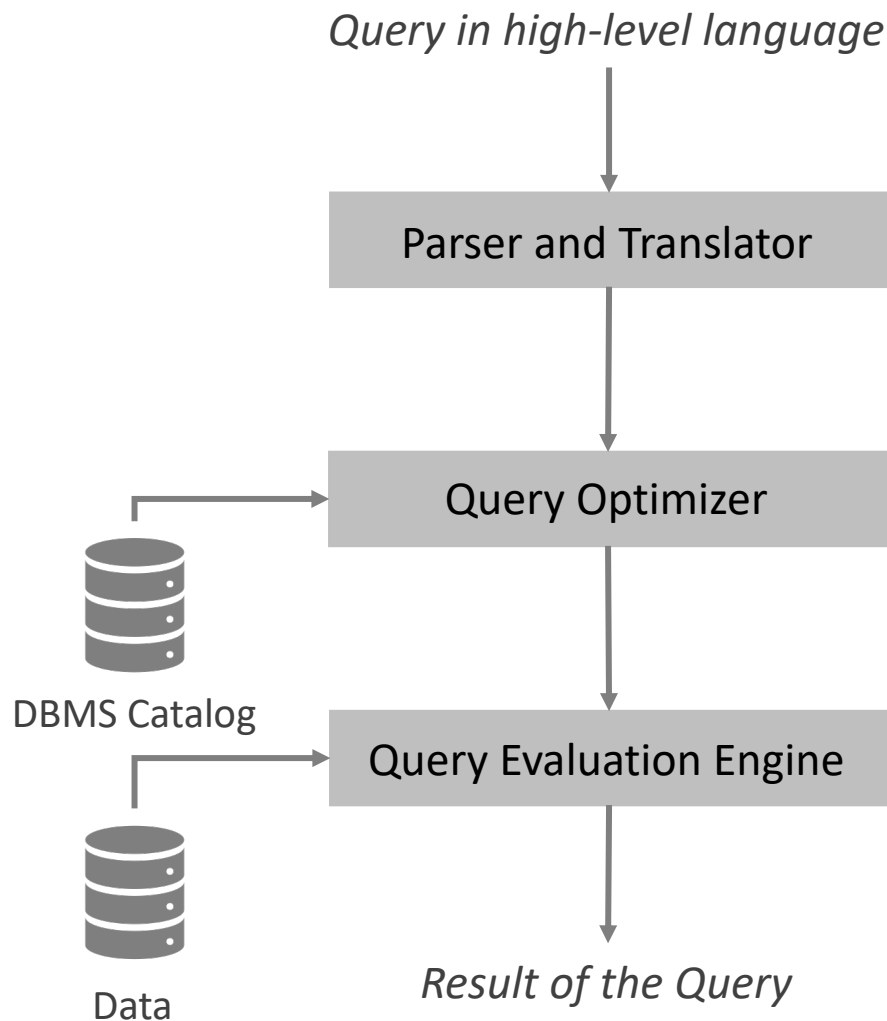
A 3-step Process

It transforms a high-level query into an equivalent and more efficient lower-level query.

Wednesday, August 1, 2018



Query Processing Example



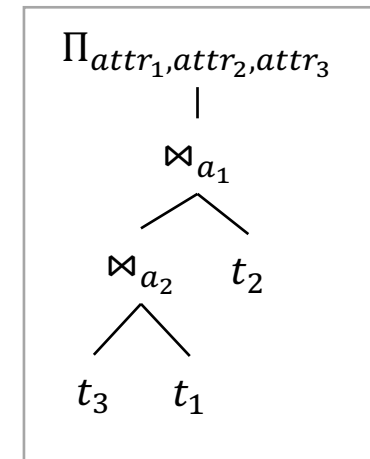
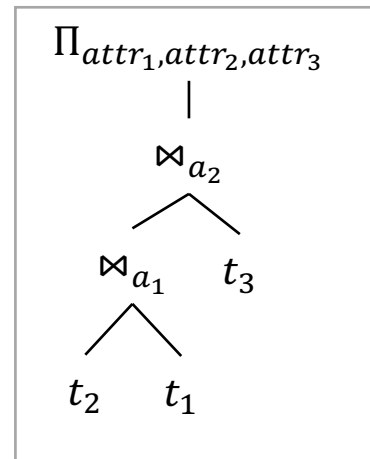
SQL expression

```
SELECT attr1, attr2, attr3
FROM t1, t2, t3
WHERE t1.a1 = t2.a1 AND t1.a2 = t3.a2
```

Algebra expression

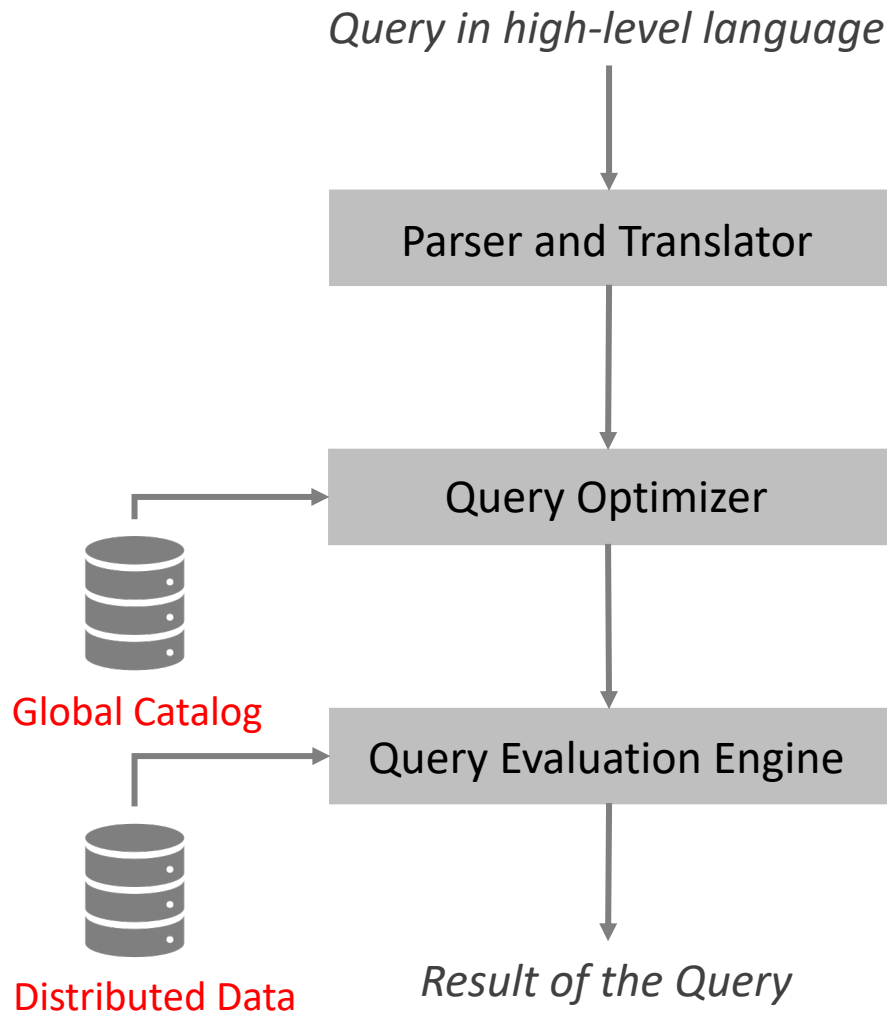
$$\Pi_{attr_1, attr_2, attr_3} (t_2 \bowtie_{s_1} t_1 \bowtie_{s_2} t_3)$$

Query execution plans



... ..

Distributed Query Processing Example



SQL expression

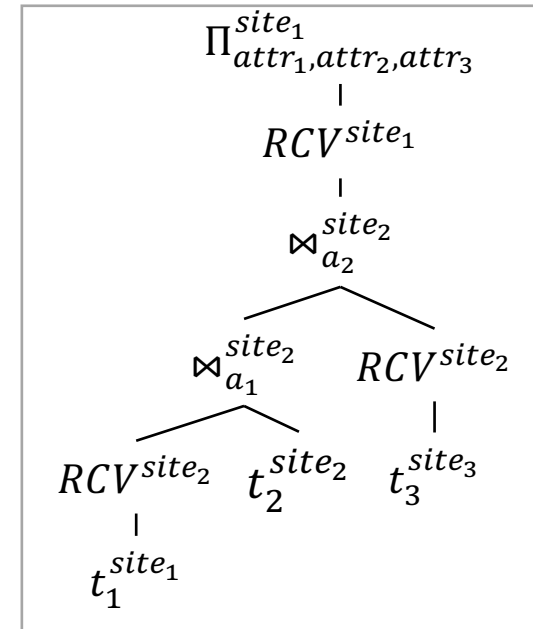
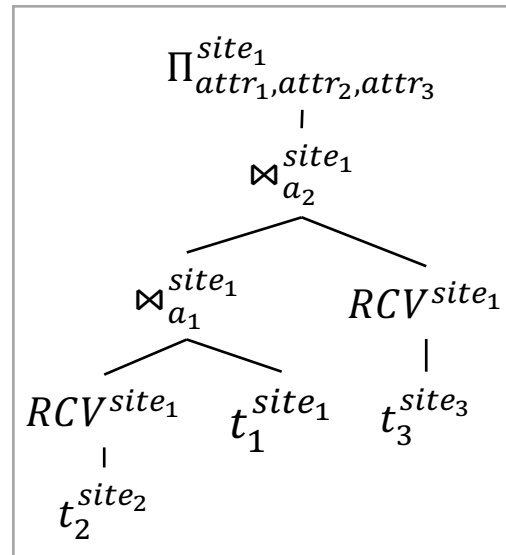
```
SELECT attr1, attr2, attr3
FROM t1, t2, t3
WHERE t1.a1 = t2.a1 AND t1.a2 = t3.a2
```

$t_1 \rightarrow \text{site}_1$
 $t_2 \rightarrow \text{site}_2$
 $t_3 \rightarrow \text{site}_3$

Algebra expression

$$\Pi_{attr_1, attr_2, attr_3} (t_2 \bowtie_{s_1} t_1 \bowtie_{s_2} t_3)$$

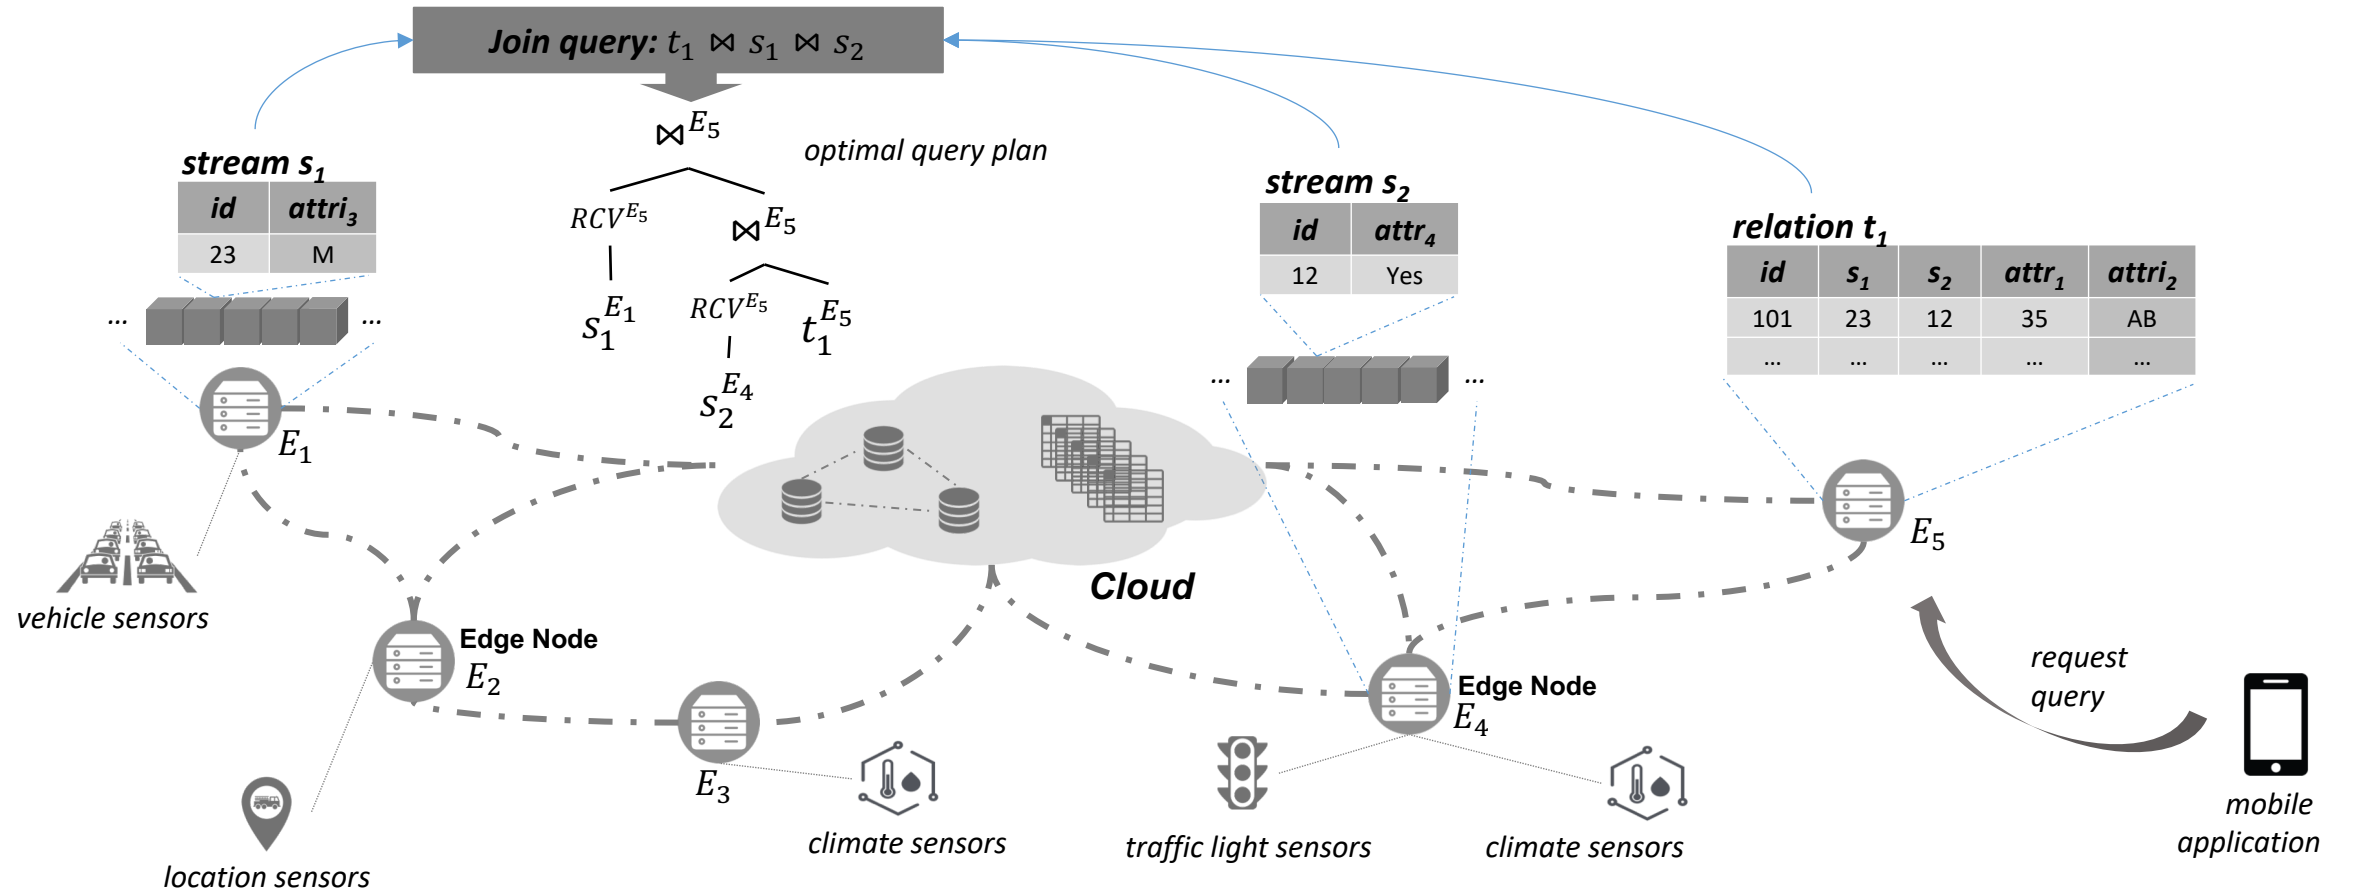
Query execution plans



Global Optimization
 +
 Local Optimization

...

Query Processing in Edge Computing



```

SELECT attr1, attr2, attr3, attr4
FROM t1, s1[RANGE 30], s2[RANGE 30]
WHERE t1.s1 = s1.id AND t1.s2 = s2.id
    
```

Challenges and Concerns : Edge vs Cloud

- Management policy
 - *Cloud servers are managed through strict and regularized policies*
 - *Edge nodes may not have the same degree of regulatory and monitoring oversight.*
 - Ship selected/projected data to edge nodes that may be untrusted/semi-trusted
 - Lead to disclosure of private information within the edge nodes.
- Latency
 - *Cloud: query in/cross data center(s) → proprietary network bandwidth*
 - Emphasis of QO is primarily on minimizing the query computation time
 - *Edge: nodes are scattered geographically with varying degrees of network connectivity*
 - A special emphasis of QO is network latency or stability

Latency Analysis

Suppose that
edge nodes are located in city A
closest cloud data center is located in city B

Edge-based approach

$$t_{edge} = \max_{i \in \{e_2, e_3\}, j \in \{(e_1, e_2), (e_1, e_3)\}} (v_i t / v_{net} + t_j) + \mathcal{T}$$

22.223 ms

Cloud-based approach

$$t_{cloud} = \max_{i \in \{e_1, e_2, e_3\}} (v_i t / v_{net} + t_{a,b}) + \mathcal{T} + \sum v_i t / v_{net} + t_{a,b}$$

84.818 ms

TABLE I
SIMULATED LOCATIONS AND THEIR NETWORK LATENCY

Location	Distance	Latency(ms)	Result
$E_1 - E_2$	d_{e_1, e_2}	$0.022 \cdot d_{e_1, e_2} + 4.862$	$t_{e_1, e_2} = 5.082$
$E_2 - E_3$	d_{e_2, e_3}	$0.022 \cdot d_{e_2, e_3} + 4.862$	$t_{e_2, e_3} = 5.202$
$E_3 - E_1$	d_{e_3, e_1}	$0.022 \cdot d_{e_3, e_1} + 4.862$	$t_{e_3, e_1} = 5.192$
$A - B$	$d_{a, b}$	$0.022 \cdot d_{a, b} + 4.862$	$t_{a, b} = 26.862$

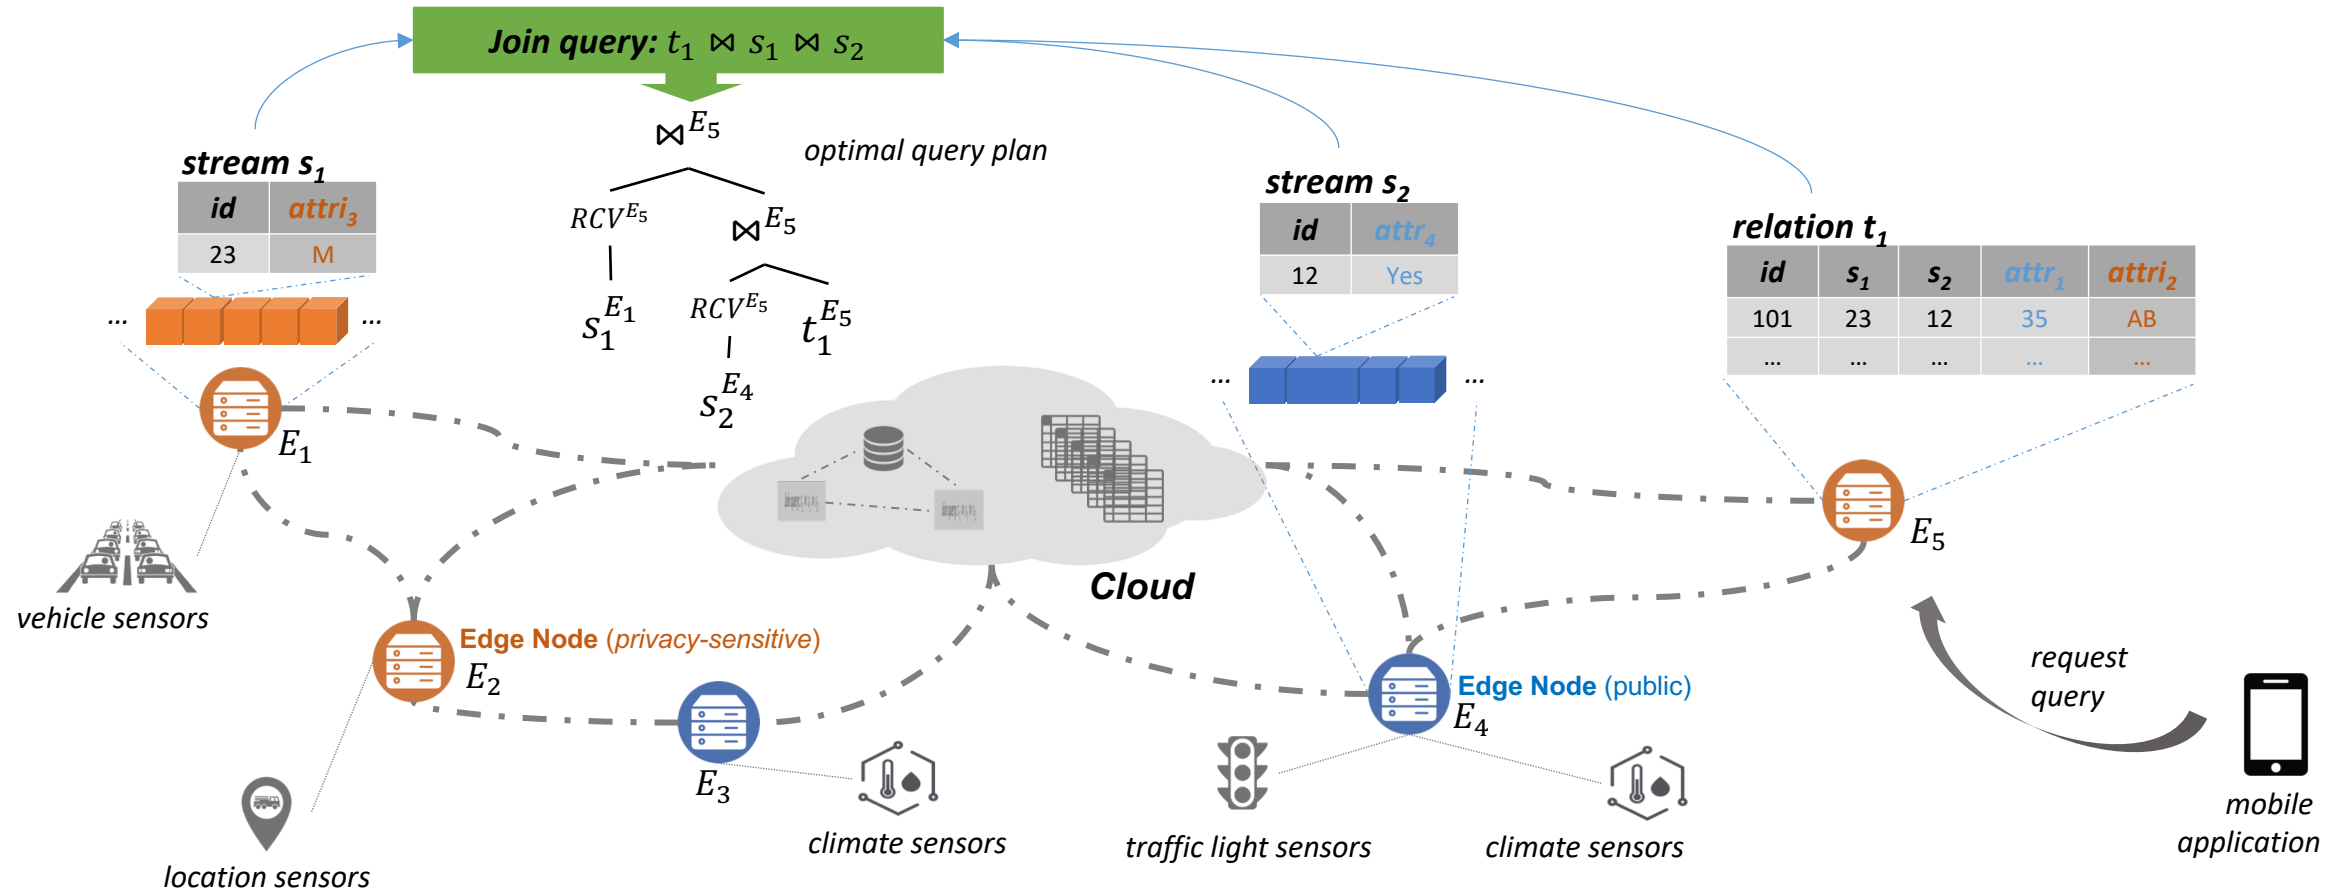
† Note that the latency of network traffic is estimated based on the distance using a linear model: $y = 0.022x + 4.862$ with coefficient of determination ($R^2 = 0.907$) proposed in [14].

‡ The distance between the data center and the city is assumed to be 1000 miles, while the distance between edge nodes is 10, 20, and 15 miles, respectively.

TABLE II
SIMULATED PARAMETER SETTINGS AND VALUES

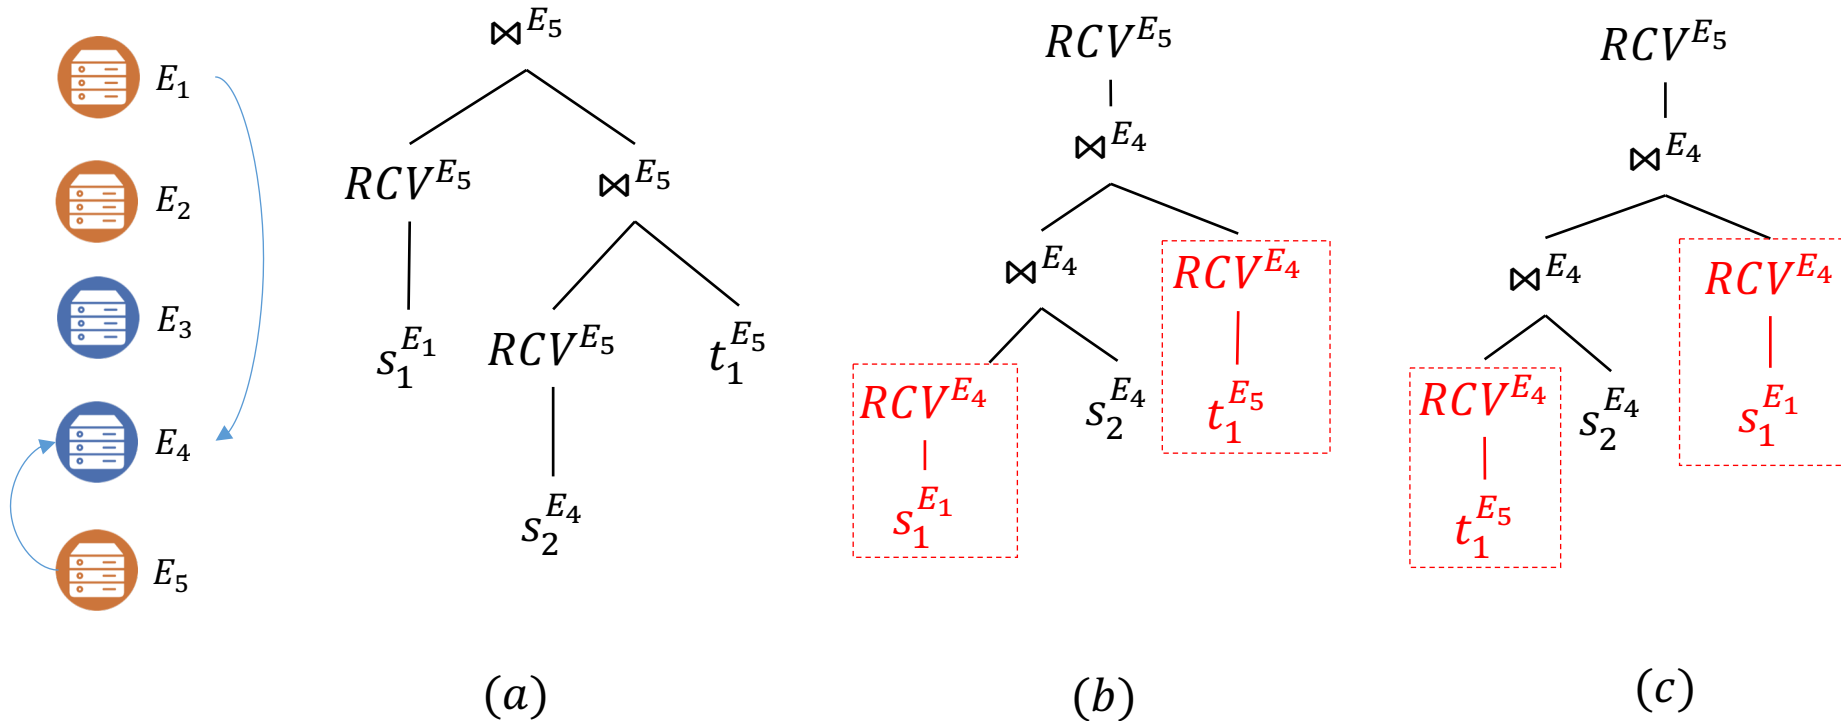
Symbol	Value	Description
t	30 min	Time interval of the query
v_{e_1}	1 KB/min	Speed of stream data generating at edge E_1
v_{e_2}	2 KB/min	Speed of stream data generating at edge E_2
v_{e_3}	3 KB/min	Speed of stream data generating at edge E_3
v_{net}	100 Mbit/s	Ethernet speed
\mathcal{T}	10 ms	Query time in a single machine

Query Processing in Edge Computing



Privacy Disclosure Risk

Suppose that E_4 is controlled by the adversary who tries to collect users' private information



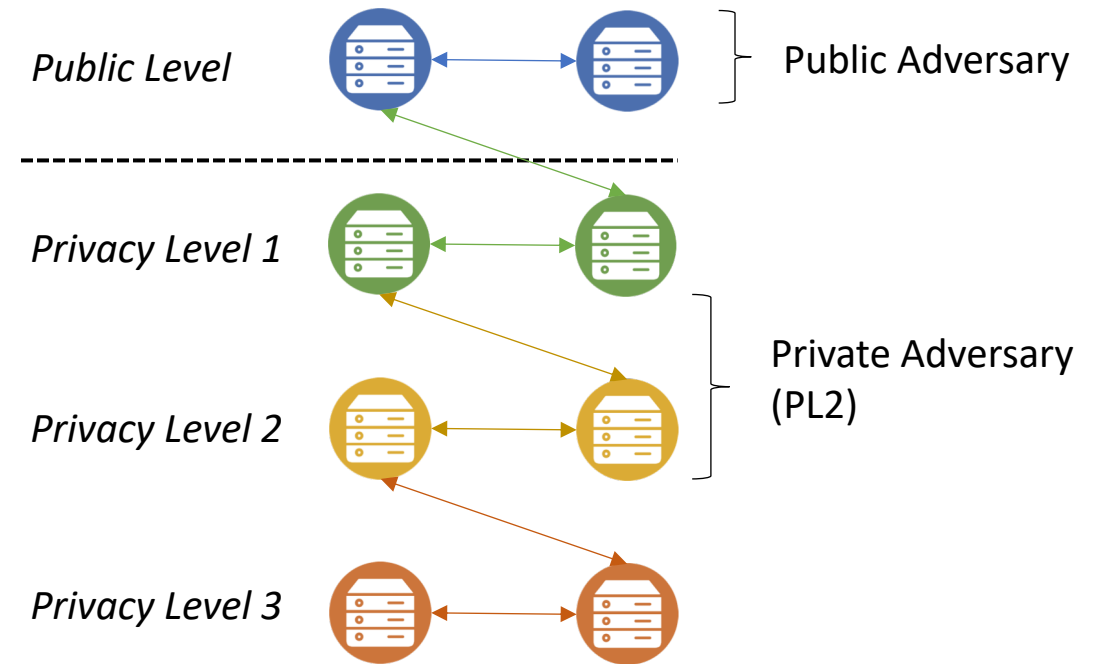
As a result, the adversary at the public edge node E_4 can acquire *the intermediate sensitive data* even if it does not have access to edge nodes where the sensitive data is stored.

Samples of query execution plan candidates

Adversary Model

- **Public Adversary**
 - *has complete control of public edge nodes*
 - *can access any data stored in public edge nodes*
- **Private Adversary**
 - *can access the private edge nodes belonging to a specific privacy level*

the adversary can access any intermediate data shipped to its controlled edge nodes during the query plan execution phase



→ the intermediate data inference attack.

Privacy Guarantee

- No privacy-sensitive information is disclosed in the query processing phase in the edge computing.
 - *if an adversary controls a public edge node*
 - it will not infer any privacy-sensitive information from monitoring the query operations
 - *even if the adversary controls a private edge node with privacy level p*
 - it cannot infer any sensitive information with privacy level higher than p

Query Guard Framework

- A traditional dynamic programming enumeration skeleton
 - *the optimal plan is generated by joining optimal sub-plans in a bottom-up manner*
- Specifically
 - *Iterative dynamic programming approach*
 - *Heuristic-based methods*

Algorithm 1: Pseudocode for QueryGuard framework

Input: A set of relations or streams $R = \{R_i\}$ with size n generated from a query Q

Output: The optimized query plan

```
1 for  $i = 1$  to  $n$  do
2   plans( $\{R_i\}$ ) := access-plans( $\{R_i\}$ )
3   LATENCY-AWARE-PRUNE(plans( $\{R_i\}$ ))
4  $toDo := R$ 
5 while  $|toDo| > 1$  do
6    $b :=$  balanced-parameter( $|toDo|, k$ )
7   for  $i = 2$  to  $b$  do
8     for all  $S \subset R$  and  $|S| = i$  do
9       plans( $S$ ) :=  $\emptyset$ 
10      for all  $O \subset S$  and  $O \neq \emptyset$  do
11        plans( $S$ ) := plans( $S$ )  $\cup$  PRIVACY-JOIN(plans( $O$ ),
12          plans( $S \setminus O$ ))
13        LATENCY-AWARE-PRUNE(plans( $S$ ))
14      find  $P, V$  with  $P \in$  plans( $V$ ),  $V \subset toDo$ ,  $|V| = k$  such that
15        eval( $P$ ) = min{eval( $P'$ ) |  $P' \in$  plans( $W$ ),  $W \subset toDo$ ,  $|W| = k$ }
16      generate new symbol:  $\mathcal{T}$ , plans( $\mathcal{T}$ ) =  $\{P\}$ 
17       $toDo = toDo - V \cup \{\mathcal{T}\}$ 
18      for all  $O \subset V$  do
19        delete(plans( $O$ ))
20 finalize-plans(plans( $R$ ))
21 LATENCY-AWARE-PRUNE(plans( $R$ ))
22 return plans( $R$ )
```

Privacy Join

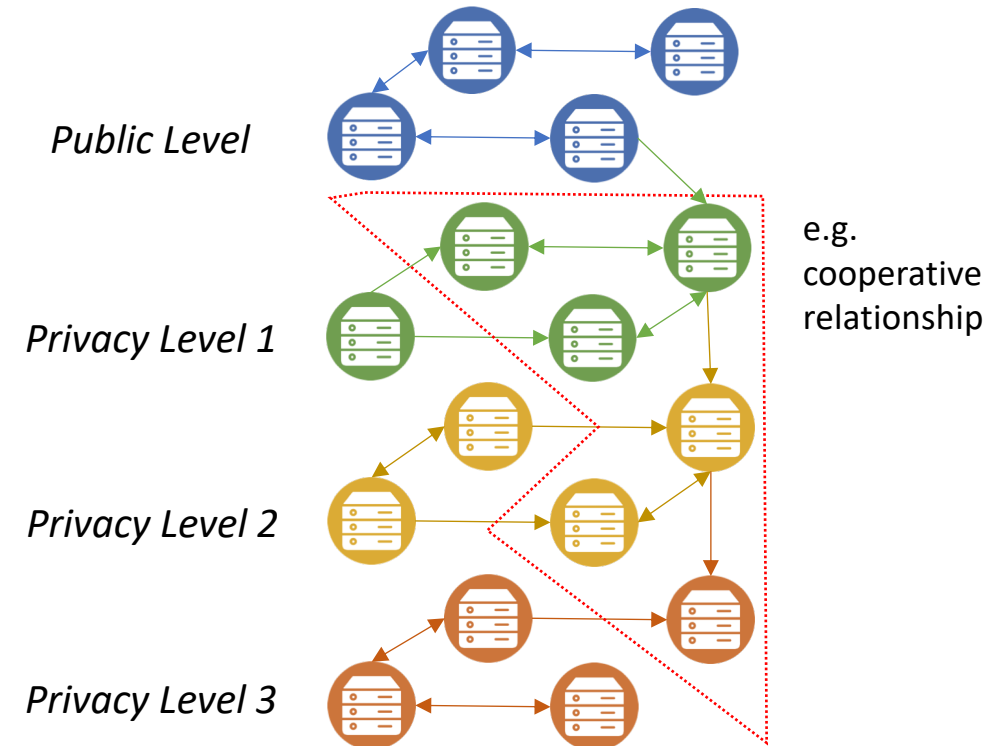
- Privacy Settings

- *Privacy Preference*

- the data is assigned a privacy preference parameter by data owner to control the data shipment scope
 - no ship out-of-scope in join operation

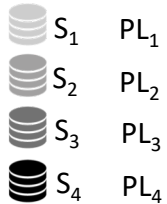
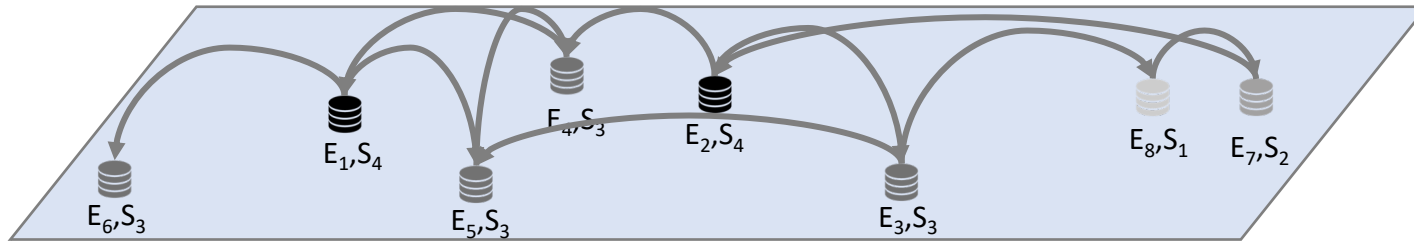
- *Privacy Level*

- each edge node is assigned a privacy level
 - the privacy level of data can be directly inferred from the privacy levels of edge nodes
 - no ship down in join operation



An illustration of the critical phases in Query Guard

Possible Joins



Privacy Levels: $PL_4 > PL_3 > PL_2 > PL_1$

Algorithm 1: Pseudocode for QueryGuard framework

Input: A set of relations or streams $R = \{R_i\}$ with size n generated from a query Q

Output: The optimized query plan

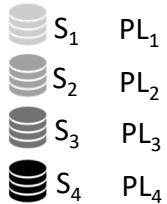
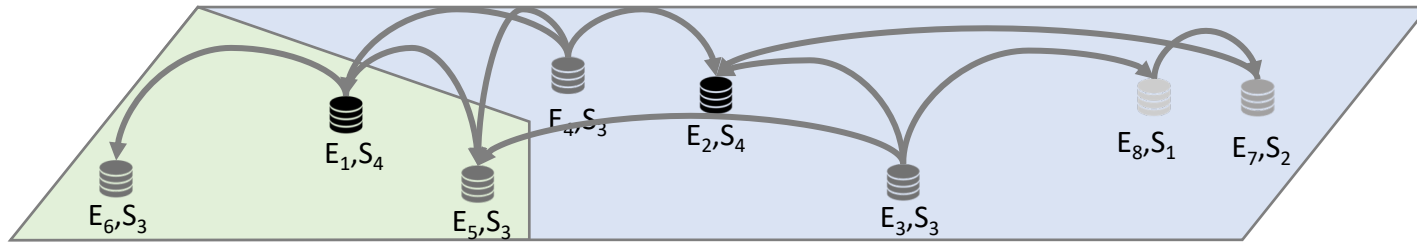
```

1 for  $i = 1$  to  $n$  do
2    $\text{plans}(\{R_i\}) := \text{access-plans}(\{R_i\})$ 
3    $\text{LATENCY-AWARE-PRUNE}(\text{plans}(\{R_i\}))$ 
4  $toDo := R$ 
5 while  $|toDo| > 1$  do
6    $b := \text{balanced-parameter}(|toDo|, k)$ 
7   for  $i = 2$  to  $b$  do
8     for all  $S \subset R$  and  $|S| = i$  do
9        $\text{plans}(S) := \emptyset$ 
10      for all  $O \subset S$  and  $O \neq \emptyset$  do
11         $\text{plans}(S) := \text{plans}(S) \cup \text{PRIVACY-JOIN}(\text{plans}(O),$ 
12           $\text{plans}(S \setminus O))$ 
13         $\text{LATENCY-AWARE-PRUNE}(\text{plans}(S))$ 
14      find  $P, V$  with  $P \in \text{plans}(V)$ ,  $V \subset toDo$ ,  $|V| = k$  such that
15         $\text{eval}(P) = \min\{\text{eval}(P') \mid P' \in \text{plans}(W), W \subset toDo, |W| = k\}$ 
16      generate new symbol:  $T$ ,  $\text{plans}(T) = \{P\}$ 
17       $toDo = toDo - V \cup \{T\}$ 
18      for all  $O \subset V$  do
19        delete( $\text{plans}(O)$ )
20 finalize-plans( $\text{plans}(R)$ )
21 LATENCY-AWARE-PRUNE( $\text{plans}(R)$ )
22 return  $\text{plans}(R)$ 

```

An illustration of the critical phases in Query Guard

Privacy-preserving Joins



Privacy Levels: $PL_4 > PL_3 > PL_2 > PL_1$

Algorithm 2: Privacy-preserving join algorithm

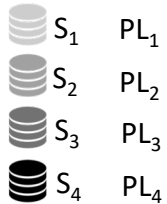
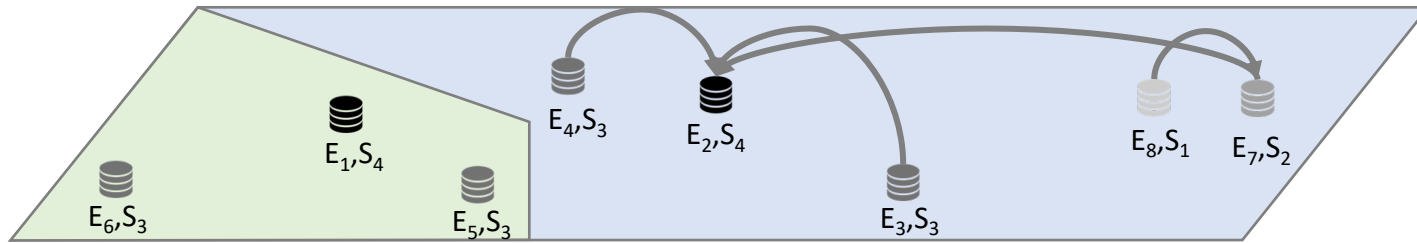
```

1 function PRIVACY-JOIN(lplans, rplans)
2   join-plans := {∅} ;
3   foreach possible edge e do
4     for plan l in lplans do
5       if PREFERENCE-CONSTRAINT(l, e) then continue;
6       lpp := LEVEL-CONSTRAINT(l, e) ;
7       if lpp.flag then continue;
8       for plan r in rplans do
9         if PREFERENCE-CONSTRAINED(r, e) then
10            continue;
11            rpp := LEVEL-CONSTRAINT(r, e) ;
12            if rpp.flag then continue;
13            join := new node(lpp.root, rpp.root, e) ;
14            join-plans.add(join) ;
15
16 return join-plans
17
18 function LEVEL-CONSTRAINT(p, e)
19   flag := false
20   if p.root.site ≠ e then
21     if p.root.P > e.P then return (true, null) ;
22     else
23       rcv_node := new node(p.root, e)
24       set P of rcv_node same to P of e.
25       p.root := rcv_node
26
27 return (flag, p)
28
29 function PREFERENCE-CONSTRAINT(p, e)
30   foreach leaf node in p do
31     λ := transmission threshold of leaf node
32     if λ is Set type then
33       if e ∉ λ then return true;
34
35 return false

```

An illustration of the critical phases in Query Guard

Latency-aware Prune



Privacy Levels: $PL_4 > PL_3 > PL_2 > PL_1$

$$f_{cost}(\mathcal{L}) = C_{cent} + \sum_{\forall (e_i, e_j) \in \mathcal{L}} (n_{bytes}^{e_i \rightarrow e_j} \cdot t_{estimate}^{e_i \rightarrow e_j})$$

$$t_{estimate}^{e_i \rightarrow e_j} = \alpha \cdot t_{avg} / n_{send}$$

$$\arctan(d_{geo}(e_i, e_j)) \cdot 2/\pi$$

Algorithm 3: Latency-aware function

```

1 function LATENCY-AWARE-PRUNE(plans(S))
2   result := {};
3   foreach site e do t[e] := null;
4   foreach plan p in plans(S) do
5     c := extract the catalog information;
6     if  $f_c(p) < f_c(t[e])$  such  $t[e] \neq null$  then  $t[e] := p$ ;
7   foreach site e do result.add(t[e]) such  $t[e] \neq null$ ;
8   return result
    
```

Experimental Evaluation

- General Setup
 - *Simulate a set of edge nodes with artificially injected network latency*
 - 15 edge nodes with specific geography information
 - Latency (ms) of the network traffic is estimated based on the distance (miles) using a linear model
 - $y = 0.022x + 4.862$

All the experiments were executed using randomly generated queries over randomly generated relations/streams that are distributed on the 15 edge nodes

EDGE NODE SIMULATION.

Edge Node Address	Privacy Level	Geography
10.0.1.{1-8}	{0,0,0,1,2,3,4,5}	Area nearby Pittsburgh, PA
10.0.1.9	0	Erie, PA
10.0.1.10	1	Philadelphia, PA
10.0.1.11	2	Allentown, PA
10.0.1.12	3	Harrisburg, PA
10.0.1.13	0	Cleveland, OH
10.0.1.14	2	Morgantown, WV
10.0.1.15	3	Washington D.C.

AN EXAMPLE OF RANDOMLY GENERATED RELATIONS/STREAMS.

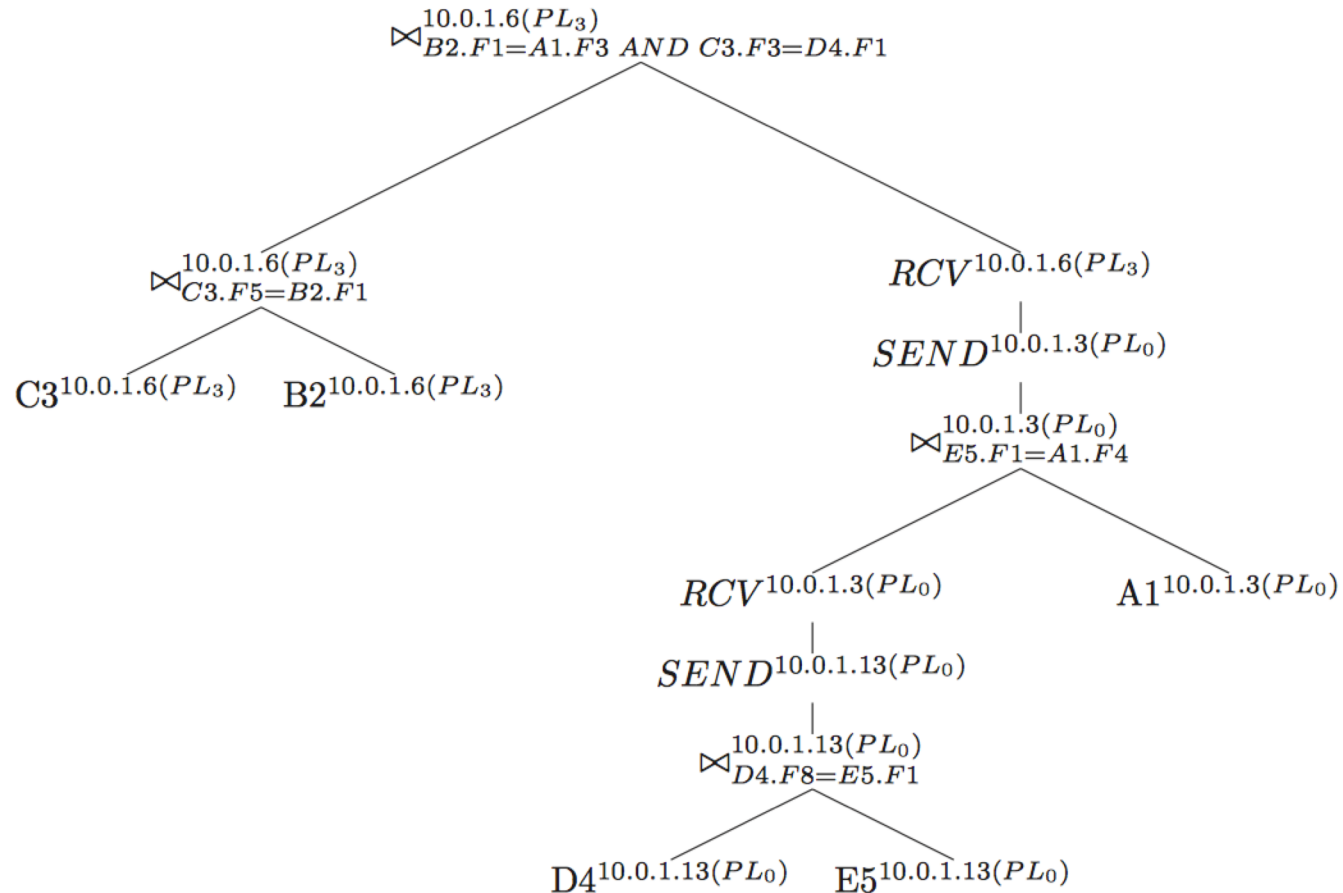
Relation/Stream	Edge Node	Transmission Threshold
A1	10.0.1.{3,4,5,7,10,14,15}	10.0.1.{1-12}
B2	10.0.1.{6,8,11,12}	10.0.1.{1-12}
C3	10.0.1.{2,6,11}	10.0.1.{1-12}
D4	10.0.1.{2,4,5,6,11,12,13}	10.0.1.{1-12}
E5	10.0.1.{4,12,13}	10.0.1.{1-12}

DISTRIBUTION OF RANDOM RELATIONS/STREAMS CARDINALITY.

Relation Type	Cardinality of Relation	Simulation Distribution
I	10-100	5%
II	100-1000	15%
III	1,000-10,000	30%
IV	10,000-100,000	30%
V	100,000-100,0000	15%
VI	1,000,000-10,000,000	5%

† The cardinality of a stream indicates the size of synopsis in DSMS.

Experimental Evaluation

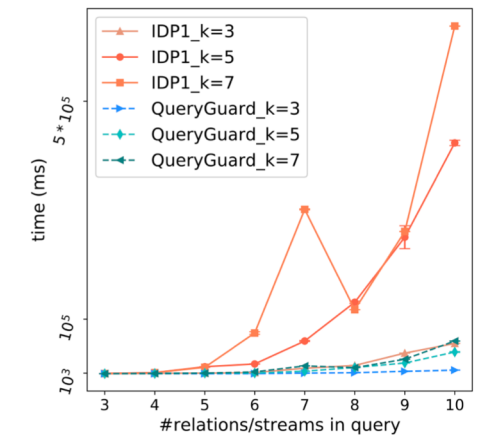
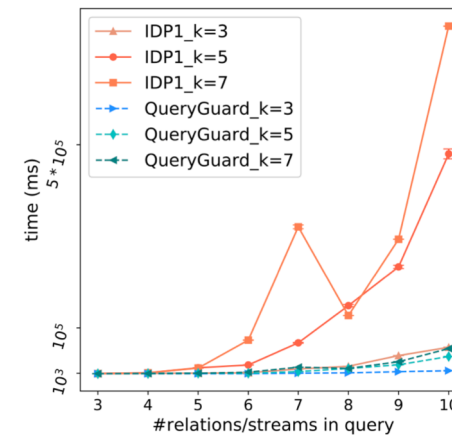
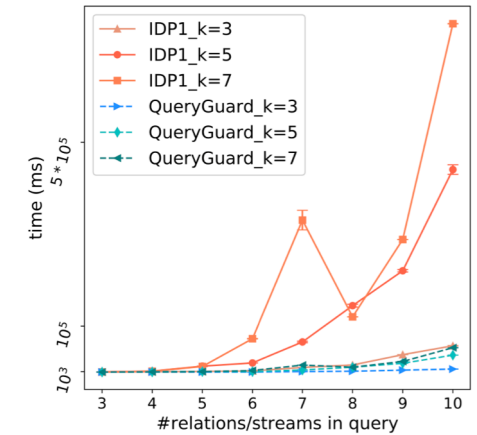
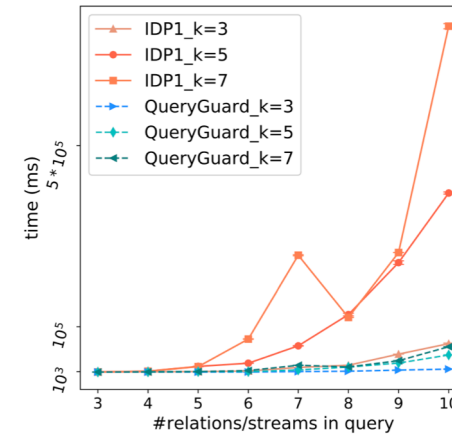


A case study of privacy-preserving processing

$A1 \boxtimes B2 \boxtimes C3 \boxtimes D4 \boxtimes E5$

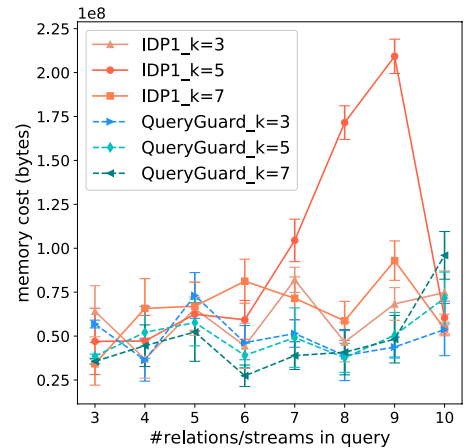
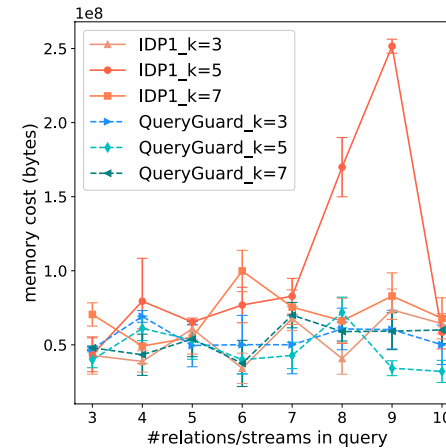
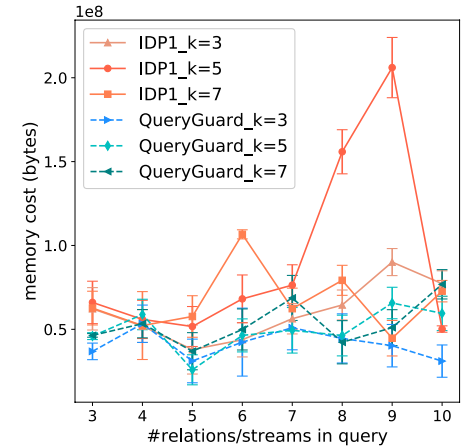
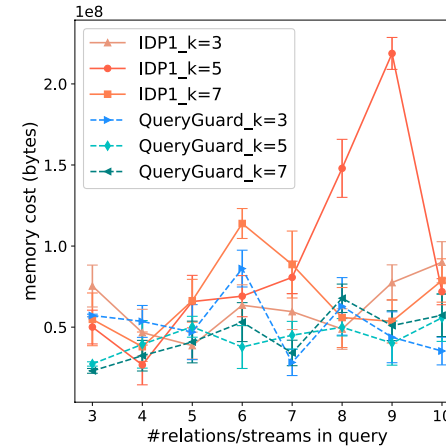
Experimental Evaluation

- Comparison to IDP1
 - *Execution Time*
 - *our proposed technique has non-negligible performance advantage in execution time*



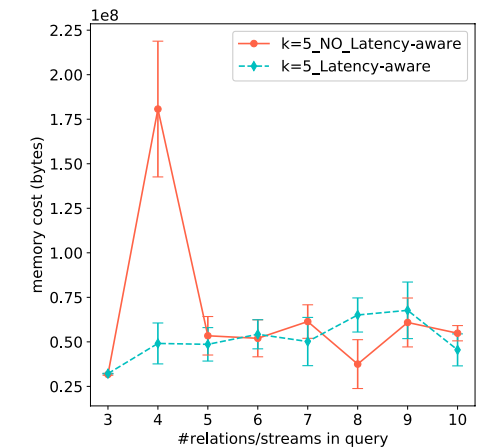
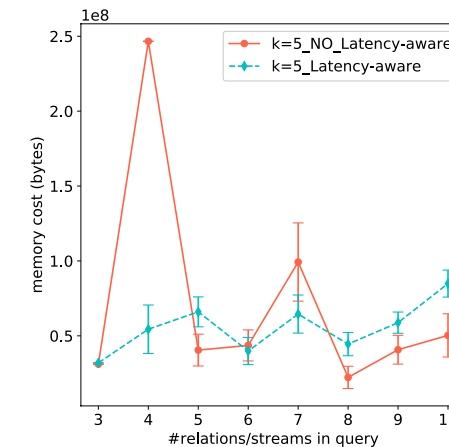
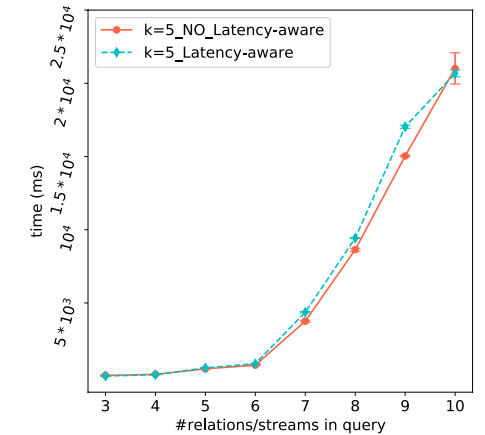
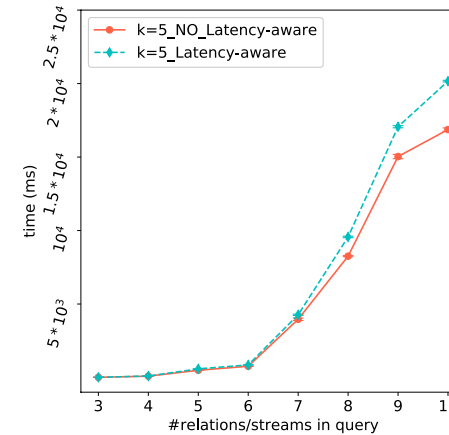
Experimental Evaluation

- Comparison to IDP1
 - *Memory Usage*
 - *our proposed technique has non-negligible performance advantage in memory usage aspects.*



Experimental Evaluation

- Effect of latency awareness setting
 - *to evaluate whether the latency-aware cost model influences the performance of our proposed framework*
- *The latency-aware setting has a negligible effect on the memory usage of the algorithm, while the execution time cost has slight growth when the relation number increase.*



Conclusion

- A privacy-preserving latency-aware query optimization framework
 - *Privacy disclosure risk analysis*
 - *Latency concerns analysis*
- *Tackled privacy-aware and latency optimized query processing in edge computing environments*
- *Evaluate the proposed techniques in terms of execution time and memory usage*
 - our results show that the proposed methods perform better than conventional techniques while achieving the intended privacy goals.

Q & A

Thanks