#### **Dissertation Prospectus**

Secure and Efficient Query Processing in Outsourced Databases Range Queries [19, 21], Point Queries [21], *k*NN Queries, **JOIN** Queries

#### Dmytro Bogatov dmytro@bu.edu

Built from 034dbe1e on January 4, 2022

Boston University Graduate School of Arts and Sciences Department of Computer Science



# INTRODUCTION AND BACKGROUND



- $\cdot$  With vast amounts of data, organizations choose to use cloud
- Challenge: solutions must be both secure and efficient
- Query types: SELECT \* FROM t1
  - Point queries: WHERE zip = '02215'
  - Range queries: WHERE age BETWEEN 18 AND 65
  - kNN queries: ORDER BY location <-> '(29.9691,-95.6972)' LIMIT 5
  - · JOIN / GROUP BY queries: INNER JOIN t2 ON (t1.k = t2.k) GROUP BY zip
- Security models for an outsourced database system
  - Snapshot adversary: steal the hard drive and RAM snapshot
  - Persistent adversary: continuously monitor the entire server



- $\cdot$  With vast amounts of data, organizations choose to use cloud
- Challenge: solutions must be both secure and efficient
- Query types: SELECT \* FROM t1
  - Point queries: WHERE zip = '02215'
  - Range queries: WHERE age BETWEEN 18 AND 65
  - kNN queries: ORDER BY location <-> '(29.9691,-95.6972)' LIMIT 5
  - · JOIN / GROUP BY queries: INNER JOIN t2 ON (t1.k = t2.k) GROUP BY zip
- Security models for an outsourced database system
  - Snapshot adversary: steal the hard drive and RAM snapshot
  - Persistent adversary: continuously monitor the entire server



- $\cdot$  With vast amounts of data, organizations choose to use cloud
- Challenge: solutions must be both secure and efficient
- Query types: SELECT \* FROM t1
  - Point queries: WHERE zip = '02215'
  - Range queries: WHERE age BETWEEN 18 AND 65
  - kNN queries: ORDER BY location <-> '(29.9691,-95.6972)' LIMIT 5
  - · JOIN / GROUP BY queries: INNER JOIN t2 ON (t1.k = t2.k) GROUP BY zip
- Security models for an outsourced database system
  - Snapshot adversary: steal the hard drive and RAM snapshot
  - Persistent adversary: continuously monitor the entire server



- $\cdot$  With vast amounts of data, organizations choose to use cloud
- Challenge: solutions must be both secure and efficient
- Query types: SELECT \* FROM t1
  - Point queries: WHERE zip = '02215'
  - Range queries: WHERE age BETWEEN 18 AND 65
  - kNN queries: ORDER BY location <-> '(29.9691,-95.6972)' LIMIT 5
  - · JOIN / GROUP BY queries: INNER JOIN t2 ON (t1.k = t2.k) GROUP BY zip
- Security models for an outsourced database system
  - Snapshot adversary: steal the hard drive and RAM snapshot
  - Persistent adversary: continuously monitor the entire server



- $\cdot$  With vast amounts of data, organizations choose to use cloud
- Challenge: solutions must be both secure and efficient
- Query types: SELECT \* FROM t1
  - Point queries: WHERE zip = '02215' D
  - Range queries: WHERE age BETWEEN 18 AND 65
  - kNN queries: ORDER BY location <-> '(29.9691,-95.6972)' LIMIT 5
  - · JOIN / GROUP BY queries: INNER JOIN t2 ON (t1.k = t2.k) GROUP BY zip
- Security models for an outsourced database system
  - Snapshot adversary: steal the hard drive and RAM snapshot
  - Persistent adversary: continuously monitor the entire server



- $\cdot$  With vast amounts of data, organizations choose to use cloud
- Challenge: solutions must be both secure and efficient
- Query types: SELECT \* FROM t1
  - Point queries: WHERE zip = '02215'
  - Range queries: WHERE age BETWEEN 18 AND 65
  - kNN queries: ORDER BY location <-> '(29.9691,-95.6972)' LIMIT 5
  - · JOIN / GROUP BY queries: INNER JOIN t2 ON (t1.k = t2.k) GROUP BY zip
- $\cdot$  Security models for an outsourced database system
  - Snapshot adversary: steal the hard drive and RAM snapshot
  - Persistent adversary: continuously monitor the entire server



- $\cdot$  With vast amounts of data, organizations choose to use cloud
- Challenge: solutions must be both secure and efficient
- Query types: SELECT \* FROM t1
  - Point queries: WHERE zip = '02215'
  - Range queries: WHERE age BETWEEN 18 AND 65
  - kNN queries: ORDER BY location <-> '(29.9691,-95.6972)' LIMIT 5
  - · JOIN / GROUP BY queries: INNER JOIN t2 ON (t1.k = t2.k) GROUP BY zip
- $\cdot$  Security models for an outsourced database system
  - Snapshot adversary: steal the hard drive and RAM snapshot
  - Persistent adversary: continuously monitor the entire server



**Dmytro Bogatov**, George Kollios, and Leonid Reyzin. A comparative evaluation of orderrevealing encryption schemes and secure range-query protocols. *Proceedings of the VLDB Endowment*, 12(8):933–947, 2019

Model: snapshot, query type: range

**Dmytro Bogatov**, Georgios Kellaris, George Kollios, Kobbi Nissim, and Adam O'Neill. *E*psolute: Efficiently querying databases while providing differential privacy. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Se-*

Model: persistent, query type: point and range

In-progress: Private kNN queries

Model: snapshot, query type: kNN

In-progress: Oblivious **JOIN** queries



**Dmytro Bogatov**, George Kollios, and Leonid Reyzin. A comparative evaluation of orderrevealing encryption schemes and secure range-query protocols. **Proceedings of the VLDB Endowment**, 12(8):933–947, 2019

Model: snapshot, query type: range

**Dmytro Bogatov**, Georgios Kellaris, George Kollios, Kobbi Nissim, and Adam O'Neill. *E*psolute: Efficiently querying databases while providing differential privacy.

In Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security (CCS '2021), 2021

Model: persistent, query type: point and range

In-progress: Private kNN queries

Model: snapshot, query type: kNN

In-progress: Oblivious **JOIN** queries



**Dmytro Bogatov**, George Kollios, and Leonid Reyzin. A comparative evaluation of orderrevealing encryption schemes and secure range-query protocols. **Proceedings of the VLDB Endowment**, 12(8):933–947, 2019

Model: snapshot, query type: range

**Dmytro Bogatov**, Georgios Kellaris, George Kollios, Kobbi Nissim, and Adam O'Neill. *E*psolute: Efficiently querying databases while providing differential privacy. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Se*-

curity (CCS '2021), 2021

Model: persistent, query type: point and range

In-progress: Private kNN queries

Model: snapshot, query type: kNN

In-progress: Oblivious JOIN queries



**Dmytro Bogatov**, George Kollios, and Leonid Reyzin. A comparative evaluation of orderrevealing encryption schemes and secure range-query protocols. **Proceedings of the VLDB Endowment**, 12(8):933–947, 2019

Model: snapshot, query type: range

**Dmytro Bogatov**, Georgios Kellaris, George Kollios, Kobbi Nissim, and Adam O'Neill. *E*psolute: Efficiently querying databases while providing differential privacy.

In Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security (CCS '2021), 2021

Model: persistent, query type: point and range

In-progress: Private kNN queries

Model: snapshot, query type: kNN

In-progress: Oblivious JOIN queries



**Dmytro Bogatov**, George Kollios, and Leonid Reyzin. A comparative evaluation of orderrevealing encryption schemes and secure range-query protocols. **Proceedings of the VLDB Endowment**, 12(8):933–947, 2019

Model: snapshot, query type: range

**Dmytro Bogatov**, Georgios Kellaris, George Kollios, Kobbi Nissim, and Adam O'Neill. *E*psolute: Efficiently querying databases while providing differential privacy.

In Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security (CCS '2021), 2021

Model: persistent, query type: point and range

In-progress: Private kNN queries
Model: snapshot, query type: kNN
In-progress: Oblivious JOIN queries
Model: persistent, query type: JOIN

Dmytro Bogatov, Angelo De Caro, Kaoutar Elkhiyaoui, and Björn Tackmann. Anonymous transactions with revocation and auditing in hyperledger fabric. In International Conference on Cryptology and Network Security. Springer, 2021



A COMPARATIVE EVALUATION OF ORDER-REVEALING ENCRYPTION SCHEMES AND SECURE RANGE-QUERY PROTOCOLS [19]

- Model: snapshot, query type: range
- Performance / security tradeoff
- Heterogeneous security definitions and leakage profiles
- $\cdot$  Performance not well-understood
  - $\cdot$  Some schemes are not even implemented
  - Prototype implementation at best
  - Not benchmarked against one another
  - Use different primitive implementations
  - Each claims to be practical and secure

### Our solution

- Analyzed security and leakages of the constructions under **a common framework**
- · Analyzed theoretically performance of the constructions
  - Implemented and ran experiments
    - Implemented 5 OPE / ORE schemes and 5 range query protocols
    - Used same language, framework and primitive implementations
    - Benchmarked primitives execution times
    - Counted invocations of primitives and I/O requests



- Model: snapshot, query type: range
- Performance / security tradeoff
- Heterogeneous security definitions and leak age profiles
- $\cdot$  Performance not well-understood
  - $\cdot$  Some schemes are not even implemented
  - Prototype implementation at best
  - Not benchmarked against one another
  - Use different primitive implementations
  - Each claims to be practical and secure

# Our solution

- Analyzed security and leakages of the constructions under **a common framework**
- • Analyzed theoretically performance of the constructions
  - Implemented and ran experiments
    - Implemented 5 OPE / ORE schemes and 5 range query protocols
    - Used same language, framework and primitive implementations
    - Benchmarked primitives execution times
    - Counted invocations of primitives and I/O requests



Epsolute: Efficiently Querying Databases While Providing Differential Privacy [21]

- Previous solutions work in the snapshot model (adversary steals the hard drive)
- What about persistent adversary (malicious script with root permissions)?
   Model: persistent, query type: point and range
- $\cdot$  Need to protect access pattern and communication volume
- Using ORAM to hide the access pattern Expensive, each request costs  $\mathcal{O}(\log n) \rightarrow ORAM definition$
- Adding fake records (noise) to the answer to hide the result size How much noise to add to have a guarantee and the least overhead? Adding a constant or a uniformly sampled noise is not an option Differential Privacy!



- Previous solutions work in the snapshot model (adversary steals the hard drive)
- What about **persistent** adversary (malicious script with **root** permissions)? Model: **persistent**, query type: **point** and **range**
- $\cdot$  Need to protect access pattern and communication volume
- Using ORAM to hide the access pattern Expensive, each request costs  $\mathcal{O}(\log n)$  (\*) ORAM definition
- Adding fake records (noise) to the answer to hide the result size How much noise to add to have a guarantee and the least overhead? Adding a constant or a uniformly sampled noise is not an option Differential Privacy!



- Previous solutions work in the snapshot model (adversary steals the hard drive)
- What about **persistent** adversary (malicious script with **root** permissions)? Model: **persistent**, query type: **point** and **range**
- Need to protect access pattern and communication volume
- Using ORAM to hide the access pattern Expensive, each request costs  $\mathcal{O}(\log n) \cong ORAM \text{ definition}$
- Adding fake records (noise) to the answer to hide the result size How much noise to add to have a guarantee and the least overhead? Adding a constant or a uniformly sampled noise is not an option Differential Privacy!



A randomized algorithm A is  $(\epsilon, \delta)$ -differentially private if for all  $\mathcal{D}_1 \sim \mathcal{D}_2 \in \mathcal{X}^n$ , and for all subsets  $\mathcal{O}$  of the output space of A,

$$\Pr\left[\mathsf{A}\left(\mathcal{D}_{1}\right)\in\mathcal{O}\right]\leq\exp(\epsilon)\cdot\Pr\left[\mathsf{A}\left(\mathcal{D}_{2}\right)\in\mathcal{O}\right]+\delta\;.$$

### How to make sense of it?

- Differential Privacy is a property of an algorithm IVE  $\Gamma$ SITy What about  $\epsilon$  and  $\delta$ ?
- How to construct such an algorithm? Laplace Perturbation Method!
- What if negative value is sampled?

Cannot truncate one side, must shift entire distribution



A randomized algorithm A is  $(\epsilon, \delta)$ -differentially private if for all  $\mathcal{D}_1 \sim \mathcal{D}_2 \in \mathcal{X}^n$ , and for all subsets  $\mathcal{O}$  of the output space of A,

$$\Pr\left[\mathsf{A}\left(\mathcal{D}_{1}\right)\in\mathcal{O}\right]\leq\exp(\epsilon)\cdot\Pr\left[\mathsf{A}\left(\mathcal{D}_{2}\right)\in\mathcal{O}\right]+\delta\;.$$

### How to make sense of it?

- Differential Privacy is a property of an algorithm What about  $\epsilon$  and  $\delta$ ?
- How to construct such an algorithm? Laplace Perturbation Method!
- What if negative value is sampled?

Cannot truncate one side, must shift entire distribution



A randomized algorithm A is  $(\epsilon, \delta)$ -differentially private if for all  $\mathcal{D}_1 \sim \mathcal{D}_2 \in \mathcal{X}^n$ , and for all subsets  $\mathcal{O}$  of the output space of A,

$$\Pr\left[\mathsf{A}\left(\mathcal{D}_{1}\right)\in\mathcal{O}\right]\leq\exp(\epsilon)\cdot\Pr\left[\mathsf{A}\left(\mathcal{D}_{2}\right)\in\mathcal{O}\right]+\delta\;.$$

### How to make sense of it?

- Differential Privacy is a property of an algorithm What about  $\epsilon$  and  $\delta$ ?
- How to construct such an algorithm? Laplace Perturbation Method!
- What if negative value is sampled? Cannot truncate one side, must shift entire distribution



A randomized algorithm A is  $(\epsilon, \delta)$ -differentially private if for all  $\mathcal{D}_1 \sim \mathcal{D}_2 \in \mathcal{X}^n$ , and for all subsets  $\mathcal{O}$  of the output space of A,

$$\Pr\left[\mathsf{A}\left(\mathcal{D}_{1}\right)\in\mathcal{O}\right]\leq\exp(\epsilon)\cdot\Pr\left[\mathsf{A}\left(\mathcal{D}_{2}\right)\in\mathcal{O}\right]+\delta\;.$$

### How to make sense of it?

- Differential Privacy is a property of an algorithm What about  $\epsilon$  and  $\delta$ ?
- How to construct such an algorithm? Laplace Perturbation Method!
- What if negative value is sampled?

Cannot truncate one side, must shift entire distribution



A randomized algorithm A is  $(\epsilon, \delta)$ -differentially private if for all  $\mathcal{D}_1 \sim \mathcal{D}_2 \in \mathcal{X}^n$ , and for all subsets  $\mathcal{O}$  of the output space of A,

$$\Pr\left[\mathsf{A}\left(\mathcal{D}_{1}\right)\in\mathcal{O}\right]\leq\exp(\epsilon)\cdot\Pr\left[\mathsf{A}\left(\mathcal{D}_{2}\right)\in\mathcal{O}\right]+\delta\;.$$

### How to make sense of it?

- Differential Privacy is a property of an algorithm What about  $\epsilon$  and  $\delta$ ?
- How to construct such an algorithm? Laplace Perturbation Method!
- What if negative value is sampled? Cannot truncate one side, must shift entire distribution





### Definition (Computationally Differentially Private Outsourced Database System (CDP-ODB))

We say that an outsourced database system  $\Pi$  is  $(\epsilon, \delta)$ -computationally differentially private (a.k.a. CDP-ODB) if for every polynomial time distinguishing adversary  $\mathcal{A}$ , for every neighboring databases  $\mathcal{D} \sim \mathcal{D}'$ , and for every query sequence  $q_1, \ldots, q_m \in \mathcal{Q}^m$  where  $m = \text{poly}(\lambda)$ ,

$$\Pr \left[ \mathcal{A} \left( 1^{\lambda}, \mathsf{VIEW}_{\Pi, \$} \left( \mathcal{D}, q_1, \dots, q_m \right) \right) = 1 \right] \leq \\ \exp \epsilon \cdot \Pr \left[ \mathcal{A} \left( 1^{\lambda}, \mathsf{VIEW}_{\Pi, \$} \left( \mathcal{D}', q_1, \dots, q_m \right) \right) = 1 \right] + \delta + \operatorname{negl}(\lambda) ,$$

the probability is over the randomness of the distinguishing adversary  $\mathcal{A}$  and the protocol  $\Pi$ . Note:

- Entire view of the adversary is DP-protected
- Implies protection against **communication volume** and **access pattern** leakages
- Query sequence  $q_1, \ldots, q_m \in \mathcal{Q}^m$  is fixed (more on that next)
- · negl( $\lambda$ ) needed for the computational (as opposed to information-theoretical) DP definition



### Definition (Computationally Differentially Private Outsourced Database System (CDP-ODB))

We say that an outsourced database system  $\Pi$  is  $(\epsilon, \delta)$ -computationally differentially private (a.k.a. CDP-ODB) if for every polynomial time distinguishing adversary  $\mathcal{A}$ , for every neighboring databases  $\mathcal{D} \sim \mathcal{D}'$ , and for every query sequence  $q_1, \ldots, q_m \in \mathcal{Q}^m$  where  $m = \text{poly}(\lambda)$ ,

$$\Pr\left[\mathcal{A}\left(1^{\lambda}, \mathsf{View}_{\Pi, S}\left(\mathcal{D}, q_{1}, \ldots, q_{m}\right)\right) = 1\right] \leq \Pr\left[\mathcal{A}\left(1^{\lambda}, \mathsf{View}_{\Pi, S}\left(\mathcal{D}, q_{1}, \ldots, q_{m}\right)\right)\right] \leq \Pr\left[\mathcal{A}\left(1^{\lambda}, \mathsf{View}_{\Pi, S}\left(\mathcal{D}, q_{1}, \ldots, q_{m}\right)\right)\right]$$

 $\exp \epsilon \cdot \Pr \left[ \mathcal{A} \left( 1^{\lambda}, \mathsf{View}_{\Pi, \$} \left( \mathcal{D}', q_1, \dots, q_m \right) \right) = 1 \right] + \delta + \mathsf{negl}(\lambda) ,$ 

the probability is over the randomness of the distinguishing adversary  $\mathcal{A}$  and the protocol  $\Pi$ . Note:

- Entire view of the adversary is DP-protected
- Implies protection against **communication volume** and **access pattern** leakages
- Query sequence  $q_1, \ldots, q_m \in \mathcal{Q}^m$  is fixed (more on that next)
- $\cdot$  negl( $\lambda$ ) needed for the computational (as opposed to information-theoretical) DP definition



### Why is the query sequence $q_1, \ldots, q_m \in Q^m$ fixed?

- Suppose neighboring medical databases differ in one record with a rare diagnosis "Alzheimer's disease"
- A medical professional, who is **a user, not an adversary** queries the database
  - for that diagnosis first
    SELECT name FROM patients WHERE condition = 'ALZ'
  - if there is a record, she queries the senior patients next
     SELECT name FROM patients WHERE age >= 65
  - otherwise she queries the general population, resulting in many more records SELECT name FROM patients
- $\cdot$  Adversary can know the answer to the first query by observing result size of the second
- Efficient system cannot return nearly the same number of records in both cases, thus, the adversary can distinguish



### Why is the query sequence $q_1, \ldots, q_m \in Q^m$ fixed?

- Suppose neighboring medical databases differ in one record with a rare diagnosis "Alzheimer's disease"
- A medical professional, who is a user, not an adversary queries the database
  - for that diagnosis first
    SELECT name FROM patients WHERE condition = 'ALZ'
  - if there is a record, she queries the senior patients next
     SELECT name FROM patients WHERE age >= 65
  - otherwise she queries the general population, resulting in many more records SELECT name FROM patients
- Adversary can know the answer to the first query by observing result size of the second
- Efficient system cannot return nearly the same number of records in both cases, thus, the adversary can distinguish



### Why is the query sequence $q_1, \ldots, q_m \in Q^m$ fixed?

- Suppose neighboring medical databases differ in one record with a rare diagnosis "Alzheimer's disease"
- A medical professional, who is a user, not an adversary queries the database
  - for that diagnosis first
    SELECT name FROM patients WHERE condition = 'ALZ'
  - if there is a record, she queries the senior patients next
     SELECT name FROM patients WHERE age >= 65
  - otherwise she queries the general population, resulting in many more records SELECT name FROM patients
- $\cdot$  Adversary can know the answer to the first query by observing result size of the second
- Efficient system cannot return nearly the same number of records in both cases, thus, the adversary can distinguish



#### Single-Threaded *E*psolute protocol





- Single-threaded version is prohibitively slow, must parallelize Assume single-threaded solution generates r = 1500 real and f = 500 noisy records
- Split  $\mathcal{U}$  and  $\mathcal{S}$  state into *m* ORAMs, run as separate machines (assume *m* = 4)
- $\cdot$  Partition records randomly (by ID) into *m* partitions, generate *m* inverted indexes
- $\cdot$  What to do about  $\mathcal{DS}?$

No- $\gamma$  method:  $\mathcal{DS}$  per ORAM

- $\cdot$  Composition of disjoint datasets: take max  $\epsilon$
- Each ORAM incurs noise comparable to *f*
- Win by splitting ORAM work *r* into *m* partitions and lose by multiplying noise *f* times *m*
- That is, each ORAM is processing  $\frac{r}{m} + f = 875$  records in parallel

 $\gamma$ -method: shared  $\mathcal{DS}$ 

- Same number of total records per ORAM
- Generated noise is larger than f (say, 2f)
- But it is split among *m* ORAMs
- That is, each ORAM is processing  $\frac{r+2f}{m} = 625$  records in parallel



- Single-threaded version is prohibitively slow, must parallelize Assume single-threaded solution generates r = 1500 real and f = 500 noisy records
- Split  $\mathcal U$  and  $\mathcal S$  state into *m* ORAMs, run as separate machines (assume *m* = 4)
- $\cdot$  Partition records randomly (by ID) into *m* partitions, generate *m* inverted indexes
- $\cdot$  What to do about  $\mathcal{DS}?$

# No- $\gamma$ method: $\mathcal{DS}$ per ORAM

- Composition of disjoint datasets: take max  $\epsilon$
- $\cdot$  Each ORAM incurs noise comparable to f
- Win by splitting ORAM work *r* into *m* partitions and lose by multiplying noise *f* times *m*
- That is, each ORAM is processing  $\frac{r}{m} + f = 875$  records in parallel

 $\gamma$ -method: shared  $\mathcal{DS}$ 

- Same number of total records per ORAM
- Generated noise is larger than f (say, 2f)
- But it is split among *m* ORAMs
- That is, each ORAM is processing  $\frac{r+2f}{m} = 625$  records in parallel



- Single-threaded version is prohibitively slow, must parallelize Assume single-threaded solution generates r = 1500 real and f = 500 noisy records
- Split  $\mathcal{U}$  and  $\mathcal{S}$  state into *m* ORAMs, run as separate machines (assume m = 4)
- Partition records randomly (by ID) into *m* partitions, generate *m* inverted indexes
- $\cdot$  What to do about  $\mathcal{DS}?$

# No- $\gamma$ method: $\mathcal{DS}$ per ORAM

- Composition of disjoint datasets: take max  $\epsilon$
- $\cdot$  Each ORAM incurs noise comparable to f
- Win by splitting ORAM work *r* into *m* partitions and lose by multiplying noise *f* times *m*
- That is, each ORAM is processing  $\frac{r}{m} + f = 875$  records in parallel

# $\gamma\text{-method:}$ shared $\mathcal{DS}$

- Same number of total records per ORAM
- Generated noise is larger than f (say, 2f)
- $\cdot$  But it is split among *m* ORAMs
- That is, each ORAM is processing  $\frac{r+2f}{m} = 625$  records in parallel



### Parallel $\mathcal{E}$ psolute diagram (with improvements)






range 10<sup>4</sup>.





Scalability measurements for  $\Pi_{\gamma}$  (shared  $\mathcal{DS}$ ) and  $\Pi_{no-\gamma}$  ( $\mathcal{DS}$  per ORAM)



## WORK-IN-PROGRESS: PRIVATE *k*NN QUERIES

- Model: snapshot, query type: kNN in arbitrary dimensions
- Input: vector of real numbers, query: return k "closest" inputs to given vector
   Distance can be L<sub>p</sub> (usually, Euclidean, p = 2) or inner (dot) product
- Applications range from similarity search to geographical search Document is a vector of words/features/topics, query is to find *k* most similar documents Object on a map is a 2D vector, query is to find *k* nearest locations
- Approximate distance-comparison preserving encryption (DCPE) scheme on input and queries

 $\forall x, y, z \in \mathbb{X}$ : DIST(x, y) <DIST $(x, z) - \beta \implies$  DIST(f(x), f(y)) <DIST(f(x), f(z))

• Prove theoretically and observe empirically how accuracy of search and efficiency of attacks drop with higher security

( >> DCPE ( >> TREC and FAISS ( >> Intermediate results plot



- Model: snapshot, query type: *k*NN in arbitrary dimensions
- Input: vector of real numbers, query: return k "closest" inputs to given vector Distance can be L<sub>p</sub> (usually, Euclidean, p = 2) or inner (dot) product
- Applications range from similarity search to geographical search
   Document is a vector of words/features/topics, query is to find k most similar documents
   Object on a map is a 2D vector, query is to find k nearest locations
- Approximate distance-comparison preserving encryption (DCPE) scheme on input and queries

 $\forall x, y, z \in \mathbb{X}$ : DIST(x, y) <DIST $(x, z) - \beta \implies$  DIST(f(x), f(y)) <DIST(f(x), f(z))

• Prove theoretically and observe empirically how accuracy of search and efficiency of attacks drop with higher security

> DCPE > TREC and FAISS > Intermediate results plot



- Model: snapshot, query type: *k*NN in arbitrary dimensions
- Input: vector of real numbers, query: return k "closest" inputs to given vector Distance can be L<sub>p</sub> (usually, Euclidean, p = 2) or inner (dot) product
- Applications range from similarity search to geographical search Document is a vector of words/features/topics, query is to find k most similar documents
   Object on a map is a 2D vector, query is to find k nearest locations
- Approximate distance-comparison preserving encryption (DCPE) scheme on input and queries

 $\forall x, y, z \in \mathbb{X}$ : DIST $(x, y) < DIST(x, z) - \beta \implies DIST(f(x), f(y)) < DIST(f(x), f(z))$ 

• Prove theoretically and observe empirically how accuracy of search and efficiency of attacks drop with higher security

( >> DCPE ) ( >> TREC and FAISS ) ( >> Intermediate results plot



- Model: snapshot, query type: *k*NN in arbitrary dimensions
- Input: vector of real numbers, query: return k "closest" inputs to given vector Distance can be L<sub>p</sub> (usually, Euclidean, p = 2) or inner (dot) product
- Applications range from similarity search to geographical search Document is a vector of words/features/topics, query is to find k most similar documents
   Object on a map is a 2D vector, query is to find k nearest locations
- Approximate distance-comparison preserving encryption (DCPE) scheme on input and queries

 $\forall x, y, z \in \mathbb{X} : \mathsf{DIST}(x, y) < \mathsf{DIST}(x, z) - \beta \implies \mathsf{DIST}(f(x), f(y)) < \mathsf{DIST}(f(x), f(z))$ 

• Prove theoretically and observe empirically how accuracy of search and efficiency of attacks drop with higher security

( >> DCPE ) ( >> TREC and FAISS ) ( >> Intermediate results plot



- Model: snapshot, query type: *k*NN in arbitrary dimensions
- Input: vector of real numbers, query: return k "closest" inputs to given vector Distance can be L<sub>p</sub> (usually, Euclidean, p = 2) or inner (dot) product
- Applications range from similarity search to geographical search Document is a vector of words/features/topics, query is to find k most similar documents
   Object on a map is a 2D vector, query is to find k nearest locations
- Approximate distance-comparison preserving encryption (DCPE) scheme on input and queries

 $\forall x, y, z \in \mathbb{X}$ : DIST (x, y) <DIST  $(x, z) - \beta \implies$ DIST (f(x), f(y)) <DIST (f(x), f(z))

• Prove theoretically and observe empirically how accuracy of search and efficiency of attacks drop with higher security

( >> DCPE ( >> TREC and FAISS ( >> Intermediate results plot



# WORK-IN-PROGRESS: OBLIVIOUS JOINS

• Model: persistent, query type: inner equi-JOIN

- Input: two tables  $T_1$  and  $T_2$ , query: return a cross-product of  $T_1$  and  $T_2$  where  $T_1.k = T_2.k$ We may also consider SELECT JOIN: WHERE  $T_1.k = T_2.k$  AND  $T_1.a = 10$
- Challenge: produce JOIN result hiding both access pattern and result size
- Proposed solution:
  - use enclave (SGX) and oblivious primitives (sort, compaction)
  - construct index over join keys, add DP noise to it
  - $\cdot$  partition the data by keys to fit a partition in the enclave
  - consolidate sparse keys as an optimization
  - do inner join within partition

► Detailed Algorithm



- Model: persistent, query type: inner equi-JOIN
- Input: two tables  $T_1$  and  $T_2$ , query: return a cross-product of  $T_1$  and  $T_2$  where  $T_1.k = T_2.k$ We may also consider SELECT JOIN: WHERE  $T_1.k = T_2.k$  AND  $T_1.a = 10$
- Challenge: produce JOIN result hiding both access pattern and result size
- Proposed solution:
  - use enclave (SGX) and oblivious primitives (sort, compaction)
  - construct index over join keys, add DP noise to it
  - $\cdot$  partition the data by keys to fit a partition in the enclave
  - consolidate sparse keys as an optimization
  - do inner join within partition

► Detailed Algorithm



- Model: persistent, query type: inner equi-JOIN
- Input: two tables  $T_1$  and  $T_2$ , query: return a cross-product of  $T_1$  and  $T_2$  where  $T_1.k = T_2.k$ We may also consider SELECT JOIN: WHERE  $T_1.k = T_2.k$  AND  $T_1.a = 10$
- Challenge: produce JOIN result hiding both access pattern and result size
- Proposed solution:
  - use enclave (SGX) and oblivious primitives (sort, compaction)
  - construct index over join keys, add DP noise to it
  - $\cdot$  partition the data by keys to fit a partition in the enclave
  - consolidate sparse keys as an optimization
  - do inner join within partition

► Detailed Algorithm



- Model: persistent, query type: inner equi-JOIN
- Input: two tables  $T_1$  and  $T_2$ , query: return a cross-product of  $T_1$  and  $T_2$  where  $T_1.k = T_2.k$ We may also consider SELECT JOIN: WHERE  $T_1.k = T_2.k$  AND  $T_1.a = 10$
- Challenge: produce JOIN result hiding both access pattern and result size
- Proposed solution:
  - use enclave (SGX) and oblivious primitives (sort, compaction)
  - construct index over join keys, add DP noise to it
  - $\cdot$  partition the data by keys to fit a partition in the enclave
  - $\cdot$  consolidate sparse keys as an optimization
  - $\cdot$  do inner join within partition

> Detailed Algorithm



## **Dissertation Prospectus**

Secure and Efficient Query Processing in Outsourced Databases Range Queries [19, 21], Point Queries [21], *k*NN Queries, **JOIN** Queries

### Dmytro Bogatov dmytro@bu.edu

Built from 034dbe1e on January 4, 2022

Boston University Graduate School of Arts and Sciences Department of Computer Science



## References

- Alexandra Boldyreva, Nathan Chenette, Younho Lee, and Adam O'Neill. Order-preserving symmetric encryption. In Advances in Cryptology - EUROCRYPT 2009, pages 224–241. Springer Berlin Heidelberg, 2009.
- [2] David Cash, Feng-Hao Liu, Adam O'Neill, Mark Zhandry, and Cong Zhang. Parameter-hiding order revealing encryption. In Advances in Cryptology – ASIACRYPT 2018, pages 181–210. Springer International Publishing, 2018.
- [3] Nathan Chenette, Kevin Lewi, Stephen A. Weis, and David J. Wu. Practical order-revealing encryption with limited leakage. In *Fast Software Encryption*, pages 474–493. Springer Berlin Heidelberg, 2016.



- [4] Ioannis Demertzis, Stavros Papadopoulos, Odysseas Papapetrou, Antonios Deligiannakis, and Minos Garofalakis. Practical private range search revisited. In *Proceedings of the 2016 International Conference on Management of Data*, pages 185–198, 2016.
- [5] Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 486–503. Springer, 2006.
- [6] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*, pages 265–284. Springer, 2006.
- [7] Georg Fuchsbauer, Riddhi Ghosal, Nathan Hauke, and Adam O'Neill. Approximate distancecomparison-preserving symmetric encryption. Cryptology ePrint Archive, Report 2021/1666, 2021. https://ia.cr/2021/1666.



- [8] Oded Goldreich. Towards a theory of software protection and simulation by oblivious rams. In Proceedings of the nineteenth annual ACM symposium on Theory of computing, pages 182–194, 1987.
- [9] Oded Goldreich and Rafail Ostrovsky. Software protection and simulation on oblivious rams. *Journal of the ACM (JACM)*, 43(3):431–473, 1996.
- [10] Jeff Johnson, Matthijs Douze, and Hervé Jégou. Billion-scale similarity search with gpus. *arXiv preprint arXiv:1702.08734*, 2017.
- [11] Florian Kerschbaum. Frequency-hiding order-preserving encryption. In Proceedings of the 22Nd ACM SIGSAC Conference on Computer and Communications Security, pages 656–667. ACM, 2015.
- [12] Florian Kerschbaum and Anselme Tueno. An efficiently searchable encrypted data structure for range queries. In *Computer Security – ESORICS 2019*, pages 344–364. Springer International Publishing, 2019.



- [13] Kevin Lewi and David J. Wu. Order-revealing encryption: New constructions, applications, and lower bounds. pages 1167–1178. ACM, 2016.
- [14] Daniel S. Roche, Daniel Apon, Seung Geol Choi, and Arkady Yerukhimovich. POPE: Partial order preserving encoding. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 1131–1142. ACM, 2016.
- [15] Cetin Sahin, Victor Zakhary, Amr El Abbadi, Huijia Lin, and Stefano Tessaro. Taostore: Overcoming asynchronicity in oblivious data storage. In 2016 IEEE Symposium on Security and Privacy (SP), pages 198–217. IEEE, 2016.
- [16] Elaine Shi, T-H Hubert Chan, Emil Stefanov, and Mingfei Li. Oblivious ram with o(log<sup>3</sup> n) worst-case cost. In International Conference on The Theory and Application of Cryptology and Information Security, pages 197–214. Springer, 2011.
- [17] Emil Stefanov, Elaine Shi, and Dawn Xiaodong Song. Towards practical oblivious RAM. In *Network and Distributed System Security Symposium (NDSS)*, 2012.



- [18] Emil Stefanov, Marten van Dijk, Elaine Shi, Christopher Fletcher, Ling Ren, Xiangyao Yu, and Srinivas Devadas. Path oram: An extremely simple oblivious ram protocol. In Proceedings of the 2013 ACM SIGSAC Conference on Computer Communications Security, pages 299–310. ACM, 2013.
- [19] **Dmytro Bogatov**, George Kollios, and Leonid Reyzin. A comparative evaluation of orderrevealing encryption schemes and secure range-query protocols. *Proceedings of the VLDB Endowment*, 12(8):933–947, 2019.
- [20] **Dmytro Bogatov**, Angelo De Caro, Kaoutar Elkhiyaoui, and Björn Tackmann. Anonymous transactions with revocation and auditing in hyperledger fabric. In *International Conference on Cryptology and Network Security*. Springer, 2021.
- [21] **Dmytro Bogatov**, Georgios Kellaris, George Kollios, Kobbi Nissim, and Adam O'Neill. Epsolute: Efficiently querying databases while providing differential privacy. In Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security (CCS '2021), 2021.



[22] Dong Xie, Guanru Li, Bin Yao, Xuan Wei, Xiaokui Xiao, Yunjun Gao, and Minyi Guo. Practical private shortest path computation based on oblivious storage. In 2016 IEEE 32nd International Conference on Data Engineering (ICDE), pages 361–372. IEEE, 2016.



# APPENDIX

Boston University

Schomo	Primitive usage		Ciphertext size,	Leakage	
Scheme	Encryption	Comparison	or state size	(in addition to inherent total order)	
BCLO [1]	n HG	none	2n	pprox Top half of the bits	
CLWW [3]	n PRF	none	2n	Most-significant differing bit	
Lewi-Wu [13]	$\frac{2n/d \text{ PRP}}{2\frac{n}{d} (2^d + 1) \text{ PRF}}$ $\frac{n}{d} 2^d \text{ Hash}$	<u>n</u> Hash	$\frac{n}{d}\left(\lambda+n+2^{d+1}\right)+\lambda$	Most-significant differing block	
CLOZ [2]	n PRF n PPH 1 PRP	n <sup>2</sup> PPH	n · h	Equality pattern of most-significant differing bit	
FH-OPE [11]	1 Traversal	3 Traversals	$3 \cdot n \cdot N$	Insertion order	
▲ Back to ORE					



Drotocol	I/O requests		Lookago	Communication (result excluded)	
PTULUCUL	Construction	Query	Leakage	Construction	Query
B+ tree with ORE	$\log_B \frac{N}{B}$	$\log_B \frac{N}{B} + \frac{r}{B}$	Same as ORE	1	1
Kerschbaum [ <mark>12</mark> ]	N B	$\log_2 \frac{N}{B} + \frac{r}{B}$	Total order	$\log_2 N$	$\log_2 N$
POPE [ <mark>14</mark> ] warm POPE [ <mark>14</mark> ] cold	1	$\log_L \frac{N}{B} + \frac{r}{B}$ N/B	<b>Partial order</b> Fully hiding	1	log <sub>L</sub> N N
Logarithmic-BRC [4]	D-U S	r U	Same as SSE	SIŁY	$\log_2 N$
ORAM	$\log^2 \frac{N}{B}$	$\log_2 \frac{N}{B} \left( \log_B \frac{N}{B} + \frac{r}{B} \right)$	Fully hiding (access pattern)	$\log^2 \frac{N}{B}$	$\log^2 \frac{N}{B}$

▲ Back to ORE



## One of the experimental results



Query stage number of I/O requests



▲ Back to ORE

Access pattern is a sequence of memory accesses **y**, where each access consists of the memory *location o*, read **r** or write **w** *operation* and the *data d* to be written.

Oblivious RAM (ORAM) is a mechanism that hides the accesses pattern. More formally, ORAM is a protocol between the client  $\mathcal{C}$  (who accesses) and the server  $\mathcal{S}$  (who stores), with a guarantee that the view of the server is indistinguishable for any two sequences of the same lengths.

$ \mathbf{v}_{4}  =  \mathbf{v}_{2} $	ORAM protocol			
$ \mathbf{y}'  =  \mathbf{y}_2 $	-	: Client C	Server S	
$VIEW_{\mathfrak{S}}(\mathbf{y}_1) \approx VIEW_{\mathfrak{S}}(\mathbf{y}_2)$		: $\mathbf{y} = (\mathbf{r}, i, \bot) _{i=1}^5$		
		: (client state) O	RAM (y) (server state)	
		: $\{d_1, d_2, d_3, d_4, d_5\}$		

For example: Square Root ORAM [8], Hierarchical ORAM [9], Binary-Tree ORAM [16], Interleave Buffer Shuffle Square Root ORAM [22], TP-ORAM [17], **Path-ORAM** [18] and TaORAM [15]. ORAM incurs at least logarithmic communication overhead in the number of stored records. [9]



## $\forall x, y, x \in \mathbb{X}$ : DIST (x, y) <DIST $(x, z) - \beta \implies$ DIST (f(x), f(y)) <DIST (f(x), f(z))

## • The scheme is by Riddhi Ghosal and Adam O'Neil [7]

- Key generation: sample at random length multiplier s and seeds for samplers
- **Encrypt**: take input vector  $x \in \mathbb{R}^d$ 
  - Sample nonce *n*
  - Using nonce and seeds, sample a point a on a  $\beta$ -radius d-dimensional ball
  - New vector is extended times *s* and points to *a*
- **Decrypt**: take encrypted vector  $c \in \mathbb{R}^d$  and nonce n
  - Do same steps except *shrink* times *s* and *remove* ball component

#### ■ Back to kNN



 $\forall x, y, x \in \mathbb{X} : \text{DIST}(x, y) < \text{DIST}(x, z) - \beta \implies \text{DIST}(f(x), f(y)) < \text{DIST}(f(x), f(z))$ 

- The scheme is by Riddhi Ghosal and Adam O'Neil [7]
- Key generation: sample at random length multiplier s and seeds for samplers
- **Encrypt**: take input vector  $x \in \mathbb{R}^d$ 
  - Sample nonce *n*
  - Using nonce and seeds, sample a point a on a  $\beta$ -radius d-dimensional ball
  - New vector is extended times s and points to a
- **Decrypt**: take encrypted vector  $c \in \mathbb{R}^d$  and nonce n
  - Do same steps except *shrink* times *s* and *remove* ball component

#### ■ Back to kNN



 $\forall x, y, x \in \mathbb{X} : \mathsf{DIST}(x, y) < \mathsf{DIST}(x, z) - \beta \implies \mathsf{DIST}(f(x), f(y)) < \mathsf{DIST}(f(x), f(z))$ 

- The scheme is by Riddhi Ghosal and Adam O'Neil [7]
- Key generation: sample at random length multiplier s and seeds for samplers
- **Encrypt**: take input vector  $x \in \mathbb{R}^d$ 
  - Sample nonce *n*
  - Using nonce and seeds, sample a point a on a  $\beta$ -radius d-dimensional ball
  - New vector is extended times s and points to a
- **Decrypt**: take encrypted vector  $c \in \mathbb{R}^d$  and nonce n
  - Do same steps except *shrink* times *s* and *remove* ball component

#### ■ Back to kNN



 $\forall x, y, x \in \mathbb{X} : \mathsf{DIST}(x, y) < \mathsf{DIST}(x, z) - \beta \implies \mathsf{DIST}(f(x), f(y)) < \mathsf{DIST}(f(x), f(z))$ 

- The scheme is by Riddhi Ghosal and Adam O'Neil [7]
- Key generation: sample at random length multiplier s and seeds for samplers
- **Encrypt**: take input vector  $x \in \mathbb{R}^d$ 
  - Sample nonce *n*
  - Using nonce and seeds, sample a point a on a  $\beta$ -radius d-dimensional ball
  - New vector is extended times s and points to a
- **Decrypt**: take encrypted vector  $c \in \mathbb{R}^d$  and nonce n
  - Do same steps except *shrink* times *s* and *remove* ball component

#### ∢ Back to *k*NN



## Component: TREC dataset and FAISS [10]

- Dataset is 8.8M documents represented as vectors of 768 dimensions Thanks Hamed Zamani for the dataset
- Query is a 768-dimensional vector asking for k = 1000 closest (inner product) documents
- Original document set is a Text **RE**trieval **C**onference (TREC) test collection set of documents, set of topics (questions), and corresponding set of relevance judgments (right answers)
- FAISS [10]: GPU-enabled library for efficient similarity search and clustering of dense vectors Developed and maintained by Facebook AI
- · General algorithm: for different eta
  - Encrypt dataset with  $\beta$
  - Encrypt queryset with  $\beta$
  - Run queries with FAISS
  - Generate TREC metrics (using relevance judgments)

#### ✓ Back to kNN



- Dataset is 8.8M documents represented as vectors of 768 dimensions Thanks Hamed Zamani for the dataset
- Query is a 768-dimensional vector asking for k = 1000 closest (inner product) documents
- Original document set is a Text **RE**trieval **Conference** (TREC) test collection set of documents, set of topics (questions), and corresponding set of relevance judgments (right answers)
- FAISS [10]: GPU-enabled library for efficient similarity search and clustering of dense vectors Developed and maintained by Facebook AI
- General algorithm: for different  $\beta$ 
  - Encrypt dataset with  $\beta$
  - Encrypt queryset with  $\beta$
  - Run queries with FAISS
  - Generate TREC metrics (using relevance judgments)

#### ✓ Back to kNN



- Dataset is 8.8M documents represented as vectors of 768 dimensions Thanks Hamed Zamani for the dataset
- Query is a 768-dimensional vector asking for k = 1000 closest (inner product) documents
- Original document set is a Text **RE**trieval **C**onference (TREC) test collection set of documents, set of topics (questions), and corresponding set of relevance judgments (right answers)
- FAISS [10]: GPU-enabled library for efficient similarity search and clustering of dense vectors Developed and maintained by Facebook AI
- General algorithm: for different  $\beta$ 
  - Encrypt dataset with  $\beta$
  - Encrypt queryset with  $\beta$
  - Run queries with FAISS
  - Generate TREC metrics (using relevance judgments)





- Dataset is 8.8M documents represented as vectors of 768 dimensions Thanks Hamed Zamani for the dataset
- Query is a 768-dimensional vector asking for k = 1000 closest (inner product) documents
- Original document set is a Text **RE**trieval **C**onference (TREC) test collection set of documents, set of topics (questions), and corresponding set of relevance judgments (right answers)
- FAISS [10]: GPU-enabled library for efficient similarity search and clustering of dense vectors Developed and maintained by Facebook AI
- $\cdot$  General algorithm: for different eta
  - Encrypt dataset with  $\beta$
  - Encrypt queryset with  $\beta$
  - Run queries with FAISS
  - Generate TREC metrics (using relevance judgments)





- Dataset is 8.8M documents represented as vectors of 768 dimensions Thanks Hamed Zamani for the dataset
- Query is a 768-dimensional vector asking for k = 1000 closest (inner product) documents
- Original document set is a Text **RE**trieval **C**onference (TREC) test collection set of documents, set of topics (questions), and corresponding set of relevance judgments (right answers)
- FAISS [10]: GPU-enabled library for efficient similarity search and clustering of dense vectors Developed and maintained by Facebook AI
- + General algorithm: for different  $\beta$ 
  - Encrypt dataset with  $\beta$
  - Encrypt queryset with  $\beta$
  - Run queries with FAISS
  - Generate TREC metrics (using relevance judgments)

✓ Back to kNN





TREC metrics, result set distance and difference, for running kNN search for  $\beta \in \{0, 1, \dots, 50\}$ 



• Back to *k*NN

## Oblivious JOINs detailed algorithm

- Construct list L of the form (k, n<sub>1</sub>, n̂<sub>1</sub>, n<sub>2</sub>, n̂<sub>2</sub>), with an element per distinct key plus noise k is a join key, n and n̂ are real and noisy numbers of records with that key in corresponding input table Noise sampled to a hierarchical sanitizer from a Laplacian distribution
- Client  $\mathcal{U}$  sends sorted *L* and hierarchical sanitizer over noise counts to the server *S* Similar to *E*psolute, adversary does not learn much from noisy counts
- Server S partitions L by k, so that partition size  $(\hat{n_1} + \hat{n_2})$  is bounded and uniform Resulting mapping from keys to partitions  $\mathcal{M}(k) = i$  can be proven DP
- Consolidate sparse keys: ensure that each *bin* corresponds to at least *U* real keys *Bin* is collection of tuples for which we will do cross-product join
- Obliviously move and pad each bin/partition with dummy records Within each bin the data is sorted by input tables
- For each bin, **do cartesian product**

#### Back to Oblivious Joins


- Construct list *L* of the form  $(k, n_1, \hat{n_1}, n_2, \hat{n_2})$ , with an element per distinct key plus noise *k* is a join key, *n* and  $\hat{n}$  are real and noisy numbers of records with that key in corresponding input table Noise sampled to a hierarchical sanitizer from a Laplacian distribution
- Client  $\mathcal{U}$  sends sorted *L* and hierarchical sanitizer over noise counts to the server *S* Similar to *E*psolute, adversary does not learn much from noisy counts
- Server S partitions L by k, so that partition size  $(\hat{n_1} + \hat{n_2})$  is bounded and uniform Resulting mapping from keys to partitions  $\mathcal{M}(k) = i$  can be proven DP
- Consolidate sparse keys: ensure that each *bin* corresponds to at least *U* real keys *Bin* is collection of tuples for which we will do cross-product join
- Obliviously move and pad each bin/partition with dummy records Within each bin the data is sorted by input tables
- For each bin, **do cartesian product**



- Construct list *L* of the form  $(k, n_1, \hat{n_1}, n_2, \hat{n_2})$ , with an element per distinct key plus noise *k* is a join key, *n* and  $\hat{n}$  are real and noisy numbers of records with that key in corresponding input table Noise sampled to a hierarchical sanitizer from a Laplacian distribution
- Client  $\mathcal{U}$  sends sorted *L* and hierarchical sanitizer over noise counts to the server *S* Similar to *E*psolute, adversary does not learn much from noisy counts
- Server *S* partitions *L* by *k*, so that partition size  $(\hat{n}_1 + \hat{n}_2)$  is bounded and uniform Resulting mapping from keys to partitions  $\mathcal{M}(k) = i$  can be proven DP
- Consolidate sparse keys: ensure that each *bin* corresponds to at least *U* real keys *Bin* is collection of tuples for which we will do cross-product join
- Obliviously move and pad each bin/partition with dummy records Within each bin the data is sorted by input tables
- For each bin, do cartesian product



- Construct list *L* of the form  $(k, n_1, \hat{n_1}, n_2, \hat{n_2})$ , with an element per distinct key plus noise *k* is a join key, *n* and  $\hat{n}$  are real and noisy numbers of records with that key in corresponding input table Noise sampled to a hierarchical sanitizer from a Laplacian distribution
- Client  $\mathcal{U}$  sends sorted *L* and hierarchical sanitizer over noise counts to the server *S* Similar to *E*psolute, adversary does not learn much from noisy counts
- Server *S* partitions *L* by *k*, so that partition size  $(\hat{n}_1 + \hat{n}_2)$  is bounded and uniform Resulting mapping from keys to partitions  $\mathcal{M}(k) = i$  can be proven DP
- Consolidate sparse keys: ensure that each *bin* corresponds to at least *U* real keys *Bin* is collection of tuples for which we will do cross-product join
- Obliviously move and pad each bin/partition with dummy records Within each bin the data is sorted by input tables
- For each bin, **do cartesian product**



- Construct list *L* of the form  $(k, n_1, \hat{n_1}, n_2, \hat{n_2})$ , with an element per distinct key plus noise *k* is a join key, *n* and  $\hat{n}$  are real and noisy numbers of records with that key in corresponding input table Noise sampled to a hierarchical sanitizer from a Laplacian distribution
- Client  $\mathcal{U}$  sends sorted *L* and hierarchical sanitizer over noise counts to the server *S* Similar to *E*psolute, adversary does not learn much from noisy counts
- Server S partitions L by k, so that partition size  $(\hat{n}_1 + \hat{n}_2)$  is bounded and uniform Resulting mapping from keys to partitions  $\mathcal{M}(k) = i$  can be proven DP
- Consolidate sparse keys: ensure that each *bin* corresponds to at least *U* real keys *Bin* is collection of tuples for which we will do cross-product join
- Obliviously move and pad each bin/partition with dummy records Within each bin the data is sorted by input tables
- For each bin, do cartesian product



- Construct list *L* of the form  $(k, n_1, \hat{n_1}, n_2, \hat{n_2})$ , with an element per distinct key plus noise *k* is a join key, *n* and  $\hat{n}$  are real and noisy numbers of records with that key in corresponding input table Noise sampled to a hierarchical sanitizer from a Laplacian distribution
- Client  $\mathcal{U}$  sends sorted *L* and hierarchical sanitizer over noise counts to the server *S* Similar to *E*psolute, adversary does not learn much from noisy counts
- Server *S* partitions *L* by *k*, so that partition size  $(\hat{n}_1 + \hat{n}_2)$  is bounded and uniform Resulting mapping from keys to partitions  $\mathcal{M}(k) = i$  can be proven DP
- Consolidate sparse keys: ensure that each *bin* corresponds to at least *U* real keys *Bin* is collection of tuples for which we will do cross-product join
- Obliviously move and pad each bin/partition with dummy records Within each bin the data is sorted by input tables
- For each bin, **do cartesian product**

