COURSE ANNOUNCEMENTS

- Homework 5 has been posted, due Friday 3/24 on Gradescope
- Homework 6 will be posted this week, and due Friday 3/31

ue Friday 3/24 on Gradescope week, and due Friday 3/31

Lecture 15: Protected Database Search

REVIEW: MPC

"Secure multi-party computation ... enables different participating entities in possession of private sets of data to link and aggregate their data sets for the exclusive purpose of performing a finite number of preapproved computations without transferring or otherwise revealing any private data to each other or anyone else."



MPC DEPLOYMENTS

<u>Cybernetica</u>: VAT tax audits



<u>Google</u>: Federated machine learning



<u>BU</u>: Pay equity in Boston



Partisia: Rate credit of farmers



<u>Unbound</u>: Protect cryptographic keys







15.1 From Data to Databases











OBJECTIVE: CRYPTOGRAPHICALLY PROTECTED DATABASE SEARCH

• Return whole dataset encrypted

Utility of stored data

No server protections (encrypt data at rest)

Property preserving encryption

• Symmetric searchable encryption

Multi-party computation

INTEREST BY GOVERNMENT STATISTICAL AGENCIES

bea.gov/evidence

unstats.un.org/bigdata/ task-teams/privacy

15.2 Designing MPC for Databases

MOTIVATING USE CASE: DIGITAL HEALTH ANALYTICS

What would be the right System X for this use case?

https://www.bu.edu/hic/research/focused-research-programs/continuous-analysis-of-mobile-health-data-among-medically-vulnerable-populations/ 8

(BU Medical & Hariri Institute for Computing)

SECURE COLLABORATIVE ANALYTICS

Medical Studies

Healthcare providers

Credit score agencies

Market Analyses

End-users

SECURE COLLABORATIVE ANALYTICS

Medical Studies

Healthcare providers

Credit score agencies

Privacy-preserving advertising

End-users

Requirements:

- No information leakage to untrusted entities
- No reliance on trusted resources
- Relational analytics
- Practical performance

GOAL: END-TO-END DATA PROTECTION

Data "at rest"

Advanced Encryption Standard (AES)

Transport Layers Security (TLS)

Data "in transit"

Data "in use"

Why should we protect data in use?

CHALLENGE: HOW TO REDUCE THE MPC COST?

"Running the query entirely under MPC [...] fails to scale beyond 3,000 total records..."

"Computing a function f on millions of client inputs" [...] could potentially take an **astronomical** amount of time in a full MPC."

"The primary source of the slowdown arises from their join operators that have **hundreds of** input tuples..."

¹ N. Volgushev, M. Schwarzkopf, B. Getchell, M. Varia, A. Lapets, and A. Bestavros. Conclave: secure multi-party computation on big data. EuroSys, 2019. ² J. Bater, G. Elliott, C. Eggen, S. Goel, A. N. Kho, and J. Rogers. SMCQL: secure querying for federated databases. PVLDB, 10(6):673-684, 2017. ³ H. Corrigan-Gibbs and D. Boneh. Prio: Private, Robust, and Scalable Computation of Aggregate Statistics, NSDI, 2017.

Plaintext	Secure	Slowdown
158	$253,\!894$	$1,\!609X$
165	$159,\!145$	967X
193	$8,\!195,\!317$	$43,\!337X$

TO MAKE MPC PRACTICAL WE NEED TO RETHINK THE SYSTEM STACK

- User Interfaces
- Query Engine
- Secure Protocols
- System Runtime
- Communication
 - Hardware

The Secrecy Framework

No information leakage

No trusted resources

Complex data analytics

OPENING THE MPC BLACK BOXES

¹ X. Wang, A. J. Malozemoff, and J. Katz. *EMP-toolkit: Efficient MultiParty* computation toolkit, 2016. <u>https://github.com/emp-toolkit</u>

Secrecy

Supported optimizations:

- Logical (e.g. operator reordering)
- Physical (e.g. message batching, operator fusion)
- Protocol-specific (e.g. dual sharing)

Secrecy

Secrecy

THREAT MODEL AND GUARANTEES

<u>Semi-honest model</u>

- -
- party (but cannot alter its execution)

<u>Security guarantees</u>

- Untrusted parties do not learn anything about:
 - The original data, intermediate or output result sizes
 - The data access patterns during query execution

* Adding support for maliciously secure primitives in progress

Computing parties do not deviate from the protocol ("honest but curious") - Adversary can monitor the network and can also compromise one computing

24

Arithmetic sharing: $x = x_1 + x_2 + x_3 \pmod{2^{64}}$

Arithmetic sharing: $x = x_1 + x_2 + x_3 \pmod{2^{64}}$

Arithmetic sharing: $x = x_1 + x_2 + x_3 \pmod{2^{64}}$

EXAMPLE: SECURE MULTIPLICATION IN SECRECY

EXAMPLE: SECURE MULTIPLICATION IN SECRECY

Arithmetic sharing: $x = x_1 + x_2 + x_3 \pmod{2^{64}}$

EXAMPLE: SECURE MULTIPLICATION IN SECRECY

FROM SECURE ADD & MUL TO RELATIONAL ANALYTICS





that is data-independent

- Data access patterns do not depend on the actual shares _
- No conditionals (if-then-else)
- No data reduction -

For 3-bit numbers:

"'Number a is greater than b if the first bit they differ as we go from left to right is **1** for a and **0** for b "

If $(a > b) \{...\}$



Cleartext

To prevent information leakage, the computing parties perform an identical computation

* Both a and b are secret-shared

$a: a_2 a_1 a_0 \qquad b: b_2 b_1 b_0$





that is data-independent

- Data access patterns do not depend on the actual shares
- No conditionals (if-then-else)
- No data reduction

For 3-bit numbers: $a : a_2a_1a_0$ $b : b_2b_1b_0$

Cleartext

To prevent information leakage, the computing parties perform an identical computation

* Both a and b are secret-shared

 \oplus $(a_2 \oplus b_2 \oplus 1) \land (a_1 \oplus b_1) \land a_1$

 \oplus $(a_2 \oplus b_2 \oplus 1) \land (a_1 \oplus b_1 \oplus 1) \land ((b_0 \oplus 1) \land a_0)$

Oblivious





that is data-independent

- Data access patterns do not depend on the actual shares _
- No conditionals (if-then-else)
- No data reduction ----

 $b: b_{2}b_{1}b_{0}$ $a: a_2 a_1 a_0$ "If the most significant bits are not the same, then a is greater than b when a_2 is set" \oplus $(a_2 \oplus b_2 \oplus 1) \land (a_1 \oplus b_1) \land a_1$ \oplus $(a_2 \oplus b_2 \oplus 1) \land (a_1 \oplus b_1 \oplus 1) \land ((b_0 \oplus 1) \land a_0)$ Oblivious

For 3-bit numbers: If $(a > b) \{...\}$ $\phi = a \stackrel{?}{>} b = (a_2 \oplus b_2) \land a_2$ Cleartext

To prevent information leakage, the computing parties perform an identical computation

* Both a and b are secret-shared





To prevent information leakage, the computing parties perform an identical computation that is data-independent

- Data access patterns do not depend on the actual shares _
- No conditionals (if-then-else)
- No data reduction —

 \oplus ($a_2 \oplus b_2 \oplus 1$) \wedge ($a_1 \oplus b_1$) \wedge \oplus $(a_2 \oplus b_2 \oplus 1) \land (a_1 \oplus b_1 \oplus 1) \land ((b_0 \oplus 1) \land a_0)$

For 3-bit numbers: $a : a_2a_1a_0$ $b : b_2b_1b_0$ If $(a > b) \{...\}$ $\phi = a \stackrel{?}{>} b = (a_2 \oplus b_2) \land a_2$ "Else, *a* is greater than *b* when the second most significant bits are not the same and a_1 is set"

Cleartext

* Both a and b are secret-shared

Oblivious





that is data-independent

- Data access patterns do not depend on the actual shares _
- No conditionals (if-then-else)
- No data reduction —

 $b: b_{2}b_{1}b_{0}$ $a: a_2 a_1 a_0$ "Else, a is greater than b when a_0 is set and b_0 is not set" $\oplus (a_2 \oplus b_2 \oplus 1) \land (a_1 \oplus b_1) \land a_1$ \oplus $(a_2 \oplus b_2 \oplus 1) \land (a_1 \oplus b_1 \oplus 1) \land ((b_0 \oplus 1) \land a_0)$

For 3-bit numbers: If $(a > b) \{...\}$ $\phi = a \stackrel{?}{>} b = (a_2 \oplus b_2) \land a_2$

Cleartext

To prevent information leakage, the computing parties perform an identical computation

* Both a and b are secret-shared

Oblivious





that is data-independent

- Data access patterns do not depend on the actual shares -
- No conditionals (if-then-else)
- No data reduction -

	Employee	Salary
R	Kim	2000
	Jane	1500
	Alex	4500





* All attributes are secret-shared

To prevent information leakage, the computing parties perform an identical computation

 $\sigma(Salary > 3000)$ Employee Ф Salary Kim 2000 0 R Jane 1500 0 Alex 4500



SECRECY'S CORE CONTRIBUTIONS

- Ι.
 - Amortize network I/O
 - Make secret-sharing competitive in high-latency (WAN) environments -

Redesigned MPC primitives that work directly on relations instead of individual records







User Interfaces

Query Engine

Secure Protocols

n Runtime

Communication

dware





 $a' \stackrel{?}{\geq} b \iff \cdots$



 $a \stackrel{?}{\geq} b \iff (a_3 \oplus b_3) \wedge a_3$

- $\bigoplus ((a_3 \oplus b_3) \oplus 1) \land (a_2 \oplus b_2) \land a_2$
- $\bigoplus ((a_3 \oplus b_3) \oplus 1) \land ((a_2 \oplus b_2) \oplus 1) \land (a_1 \oplus b_1) \land a_1$
- $\bigoplus ((a_3 \oplus b_3) \oplus 1) \land ((a_2 \oplus b_2) \oplus 1) \land ((a_1 \oplus b_1) \oplus 1) \land (((a_0 \oplus 1) \land b_0) \oplus 1))$

- \bigoplus (($a_3 \oplus b'_3$) \oplus 1) \land ($a_2 \oplus b'_2$) \land a_2
- $((a_3 \oplus b'_3) \oplus 1) \land ((a_2 \oplus b'_2) \oplus 1) \land (a_1 \oplus b'_1) \land a_1$ \bigoplus
- $\bigoplus ((a_3 \oplus b'_3) \oplus 1) \land ((a_2 \oplus b'_2) \oplus 1) \land ((a_1 \oplus b'_1) \oplus 1) \land (((a_0 \oplus 1) \land b'_0) \oplus 1))$











 $a' \stackrel{?}{\geq} b \iff \cdots$

44



 $a' \stackrel{?}{\geq} b \iff \cdots$





 $a' \stackrel{?}{\geq} b \iff \cdots$





 $a' \stackrel{?}{\geq} b \iff \cdots$



EFFECT OF MESSAGE BATCHING (LAN)



* Parties deployed on AWS EC2 r5.xlarge instances (us-east-2)

- Eager: Message batching disabled (one network I/O per row)
- Batched: Message batching enabled

Lower is better





SECRECY'S CORE CONTRIBUTIONS

- - Amortize network I/O
 - Make secret-sharing competitive in high-latency (WAN) environments -

2. Analytical cost model based on secure computation and communication primitives

- Operation cost (number of MPC operations)
- Synchronization cost (number of communication rounds)
- Composition cost (extra cost of composing relational operators)

Redesigned MPC primitives that work directly on relations instead of individual records





LOGICAL TRANSFORMATION RULES (CLEARTEXT DATABASES)



SELECT P.id FROM Patients as P, Clients as C WHERE P.id = C.idAND P.zip='02446'

"Find the IDs of patients who are also clients (of an insurance company) and live in Brookline"





LOGICAL TRANSFORMATION RULES (CLEARTEXT DATABASES)



SELECT P.id FROM Patients as P, Clients as C WHERE P.id = C.id AND P.zip='02446'

Pushing the selection down reduces the size of intermediate data and improves performance





CLEARTEXT OPTIMIZATIONS ARE NOT ALWAYS EFFECTIVE UNDER MPC



Pushing the selection before the JOIN does not improve JOIN's performance under MPC (since the oblivious selection does not remove any tuples from P)





OPERATOR REORDERING STILL MAKES SENSE UNDER MPC



User Interfaces

Query Engine

Secure Protocols

n Runtime

Communication

rdware



OPERATOR REORDERING STILL MAKES SENSE UNDER MPC



SELECT DISTINCT M.id FROM Medication as M, Prescribed as P WHERE M.id = P.id

> "Find the distinct medication IDs that have been prescribed to patients"





OPERATOR REORDERING STILL MAKES SENSE UNDER MPC



 $O(n^2 \log^2 n)$ operations / messages $O(\log^2 n)$ rounds $O(n^2)$ space

* Assuming the distinct operator is based on a sorting network

SELECT DISTINCT M.id FROM Medication as M, Prescribed as P WHERE M.id = P.id

 $|\mathbf{M}| = |\mathbf{P}| = n$





EXAMPLE: DISTINCT PUSH-DOWN IN SECRECY



 $O(n^2 \log^2 n)$ operations / messages $O(\log^2 n)$ rounds $O(n^2)$ space

* Assuming the distinct operator is based on a sorting network



 $O(n^2)$ operations / messages ~ 4 × fewer rounds O(n) space



EXAMPLE: DISTINCT PUSH-DOWN IN SECRECY



 $O(n^2 \log^2 n)$ operations / messages $O(\log^2 n)$ rounds $O(n^2)$ space

* Assuming the distinct operator is based on a sorting network

 $O(n^2)$ operations / messages

 $\sim 4 \times$ fewer rounds O(n) space



EFFECT OF DISTINCT PUSH-DOWN (LAN)



* Parties deployed on AWS EC2 r5.xlarge instances (us-east-2)







SELECT M.med, COUNT(*) FROM Medication as M, Patients as P WHERE M.id = P.id GROUP-BY M.med

"Count the number of patients per prescribed medication"







Applying GROUP-BY after the join will require materializing the cartesian product $M \times P$

SELECT M.med, COUNT(*) FROM Medication as M, Patients as P WHERE M.id = P.idGROUP-BY M.med

We can decompose the aggregation in two parts and push the first (and most expensive one) down







Applying GROUP-BY after the join will require materializing the cartesian product $M \times P$



cnt is the number of times each id in M matched with an id in P during the SEMI-JOIN







 $O(n^2 \log^2 n)$ operations / messages $O(n^2)$ rounds $O(n^2)$ space

* Assuming the group-by operator is based on a sorting network



 $O(n^2)$ operations / messages O(n) rounds O(n) space





EFFECT OF JOIN-AGGREGATION DECOMPOSITION (LAN)



* Parties deployed on AWS EC2 r5.xlarge instances (us-east-2)

Lower is better



SECRECY'S CORE CONTRIBUTIONS

- - Amortize network I/O
 - Make secret-sharing competitive in high-latency (WAN) environments

2. Analytical cost model based on secure computation and communication primitives

- Operation cost (number of MPC operations)
- Synchronization cost (number of communication rounds)
- Composition cost (extra cost of composing relational operators)

3. Volcano-style query processor for vectorized MPC execution

- Logical optimizations (e.g., operator reordering and decomposition)
- Physical optimizations (e.g., message batching, operator fusion)
- Protocol-specific optimizations (e.g., dual sharing)

Redesigned MPC primitives that work directly on relations instead of individual records





SECRECY OVERVIEW



Public or private cloud, federated datacenter, on-premises cluster



SECRECY OVERVIEW



Public or private cloud, federated datacenter, on-premises cluster

66

SECRECY OVERVIEW



Public or private cloud, federated datacenter, on-premises cluster



Performance on real and synthetic queries



SUMMARY OF RESULTS

- Secrecy optimizations can improve performance by orders of magnitude
- Secrecy scales to millions of input rows 2.
- 3. Secrecy outperforms state-of-the-art frameworks
- Relational operators scale well with configurable memory footprints 4.
- Secrecy can perform millions of primitive MPC operations per second 5.

J. Liagouris, V. Kalavri, M. Faisal, M. Varia. Secrecy: Secure Collaborative Analytics on Secret-shared Data. arXiv:2102.01048, 2021.


SUMMARY OF RESULTS

- Secrecy optimizations can improve performance by orders of magnitude
- Secrecy scales to millions of input rows 2.
- 3. Secrecy outperforms state-of-the-art frameworks
- 4. Relational operators scale well with configurable memory footprints
- 5. Secrecy can perform millions of primitive MPC operations per second

J. Liagouris, V. Kalavri, M. Faisal, M. Varia. Secrecy: Secure Collaborative Analytics on Secret-shared Data. arXiv:2102.01048, 2021.





* Parties deployed in three AWS regions: us-east-2 (Ohio), us-east-1 (Virginia), and us-west-1 (California)

* Reported times are for 1000 rows per input relation

* Not optimized plans use message batching too (otherwise the cost of MPC is prohibitive)







* Parties deployed in three AWS regions: us-east-2 (Ohio), us-east-1 (Virginia), and us-west-1 (California)

* Reported times are for 1000 rows per input relation

* Not optimized plans use message batching too (otherwise the cost of MPC is prohibitive)



SECRECY's SCALING BEHAVIOR (LAN)



* Parties deployed on AWS EC2 r5.xlarge instances (us-east-2)

Rec. C. Diff scales to 2 million rows in $\sim 1.2h$

```
WITH rcd AS (
   SELECT pid, time, row_no() OVER
   (PARTITION BY pid ORDER BY time)
   FROM diagnosis
   WHERE diag=cdiff)
 SELECT DISTINCT pid
 FROM rcd r1 JOIN rcd r2 ON r1.pid = r2.pid
 WHERE r2.time - r1.time >= 15 DAYS
 AND r2.time - r1.time <= 56 DAYS
 AND r2.row_no = r1.row_no + 1
```

"Find the distinct ids of patients who have been diagnosed with cdiff and have two consecutive infections between 15 and 56 days apart"

Rows





What about more complex queries?



```
select
        o_year,
        sum(case
                 when nation = '[NATION]'
                 then volume
                 else 0
        end) / sum(volume) as mkt_share
from (
         select
                 extract(year from o_orderdate) as o_year,
                 l_extendedprice * (1-l_discount) as volume,
                 n2.n name as nation
        from
                 part,
                 supplier,
                 lineitem,
                 orders,
                 customer,
                 nation n1.
                 nation n2,
                 region
         where
                 p_partkey = l_partkey
                 and s_suppkey = l_suppkey
                 and l_orderkey = o_orderkey
                 and o custkey = c custkey
                 and c_nationkey = n1.n_nationkey
                 and n1.n_regionkey = r_regionkey
                 and r_name = '[REGION]'
                 and s nationkey = n2.n nationkey
                 and o_orderdate between date '1995-01-01' and date '1996-12-31'
                 and p_type = '[TYPE]'
        ) as all nations
group by
        o_year
order by
        o_year;
```

""This query determines how the market share of a given nation within a given region has changed over two years for a given part type"

> A naive approach would require $O(n^8 \log^2 n)$ operations under MPC





```
select
        o_year,
        sum(case
                 when nation = '[NATION]'
                 then volume
                 else 0
        end) / sum(volume) as mkt_share
from (
         select
                 extract(year from o orderdate) as o year,
                 1 extendedprice * (1-1 discount) as volume,
                 n2.n name as nation
        from
                 part,
                 supplier,
                 lineitem,
                 orders,
                 customer,
                 nation n1.
                 nation n2,
                 region
         where
                 p_partkey = l_partkey
                 and s_suppkey = l_suppkey
                 and l_orderkey = o_orderkey
                 and o custkey = c custkey
                 and c nationkey = n1.n nationkey
                 and n1.n_regionkey = r_regionkey
                 and r_name = '[REGION]'
                 and s nationkey = n2.n nationkey
                 and o_orderdate between date '1995-01-01' and date '1996-12-31'
                 and p_type = '[TYPE]'
        ) as all nations
group by
         o_year
order by
        o_year;
```

""This query determines how the market share of a given nation within a given region has changed over two years for a given part type"







```
select
        o_year,
        sum(case
                 when nation = '[NATION]'
                 then volume
                 else 0
        end) / sum(volume) as mkt_share
from (
         select
                 extract(year from o_orderdate) as o_year,
                 1_extendedprice * (1-l_discount) as volume,
                 n2.n name as nation
        from
                 part,
                 supplier,
                 lineitem,
                 orders,
                 customer.
                 nation n1.
                 nation n2,
                 region
         where
                 p_partkey = l_partkey
                 and s_suppkey = l_suppkey
                 and l_orderkey = o_orderkey
                 and o_custkey = c_custkey
                 and c_nationkey = n1.n_nationkey
                 and n1.n regionkey = r regionkey
                 and r_name = '[REGION]'
                 and s nationkey = n2.n nationkey
                 and o_orderdate between date '1995-01-01' and date '1996-12-31'
                 and p_type = '[TYPE]'
        ) as all nations
group by
         o_year
order by
        o_year;
```

""This query determines how the market share of a given nation within a given region has changed over two years for a given part type"







LOGICAL + SYSTEM OPTIMIZATIONS TO SCALE COMPLEX QUERIES



Interfaces

Query Engine

e Protocols

System Runtime

nunication

rdware





Customer

SELECT O.Oid FROM Customer as C, Orders as O WHERE $C.C_{id} = O.C_{id}$ AND C.Name=John

* All attributes are secret-shared







Customer

SELECT O.Oid FROM Customer as C, Orders as O WHERE $C.C_{id} = O.C_{id}$ AND C.Name=John

* All attributes are secret-shared

R_{i}	d	Name	C_{id}	O_{id}
С		Mary	5	
С		Ann	12	
С		Bob	189	
С		John	7	
С		Tom	66	
0			5	12
0			7	123
0			66	33
0			7	4

 $S \cup T$





SELECT O.Oid FROM Customer as C, Orders as O



* All attributes are secret-shared

R_{id}	Name	C_{id}	<i>O</i> _{id}
С	Mary	5	
С	Ann	12	
С	Bob	189	
С	John	7	
С	Tom	66	
0		5	12
0		7	123
0		66	33
0		7	4

SUT

* All attributes are secret-shared

R _{id}		Name		C _{id}	O _{id}
С		Mary		5	
С		Ann		12	
С		Bob		189	
С		John		7	
С		Tom		66	
Ο				5	12
Ο				7	123
Ο				66	33
0				7	4
$S \cup T$					

Sort on C_{id}, R_{id}

 $O(n \log^2 n)$ operations $O(\log^2 n)$ rounds $n = |S \cup T|$ O(n) space

* Assuming the distinct operator is based on a sorting network



* All attributes are secret-shared

R _{id}	Name	C_{id}	O _{id}
С	Mary	5	
С	Ann	12	
С	Bob	189	
С	John	7	
С	Tom	66	
0		5	12
0		7	123
Ο		66	33
Ο		7	4
SUT			

R _{id}	Name	C_{id}	O _{id}
Ο		5	12
С	Mary	5	
Ο		7	123
0		7	4
С	John	7	
С	Ann	12	
0		66	33
С	Tom	66	
С	Bob	189	
·	\sim		

 $S \cup T$

* All attributes are secret-shared

 O_{id}

R _{id}	Name	C_{id}	O _{id}
С	Mary	5	
С	Ann	12	
С	Bob	189	
С	John	7	
С	Tom	66	
Ο		5	12
0		7	123
0		66	33
0		7	4

 $S \cup T$



12 Apply odd-even aggregation phase to find "matched pairs of tuples" and mask the rest 123 4 $O(n \log n)$ operations $O(\log n)$ rounds $n = |S \cup T|$ O(n) space 33

Using this simple idea, we can evaluate the whole query in:

- $O(n \log^2 n)$ operations
- $O(\log^2 n)$ rounds
- O(n) space

where n is the total number of input rows across all input relations

"This query determines how the market share of a given nation within a given region has changed over two years for a given part type"







The optimization applies to all MPC queries used in existing systems







PARALLEL OBLIVIOUS SORT ON SECRECY (LAN)



* Reported times are for 2M input rows

* Parties deployed on AWS EC2 r5.xlarge instances (us-east-2)



Ongoing and Future Work



REAL-WORLD SECRECY USE CASES

Digital Health Analytics

(BU Medical & Hariri Institute for Computing)



https://www.bu.edu/hic/research/focused-research-programs/continuous-analysisof-mobile-health-data-among-medically-vulnerable-populations/

Secure cross-site analytics on OpenShift logs (BU RedHat Collaboratory)



https://www.bu.edu/rhcollab/projects/security-privacy/secure-cross-site-analytics-onopenshift-logs/



BROADER VISION A general-purpose framework for private data analysis in untrusted clouds



- + Relational analytics
- + ML workloads
- + Fully Homomorphic Encryption primitives
- + Differential Privacy
- Hardware acceleration
 for secure computation



SECRECY SUMMARY



Up to 1000x speedups for real and synthetic queries



J. Liagouris, V. Kalavri, M. Faisal, M. Varia. Secrecy: Secure Collaborative Analytics on Secret-shared Data. arXiv:2102.01048, 2021.



Millions of input rows entirely under MPC





