



# HEDAS: Secure and Efficient Distributed OLAP using Fully Homomorphic Encryption

Yu Tian, Tianxiang Shen, Qi Hu, Wei Chen, Heming Cui, and Ji  
Qi\*

[tianyuk@cs.hku.hk](mailto:tianyuk@cs.hku.hk)

Sep. 19, 2024

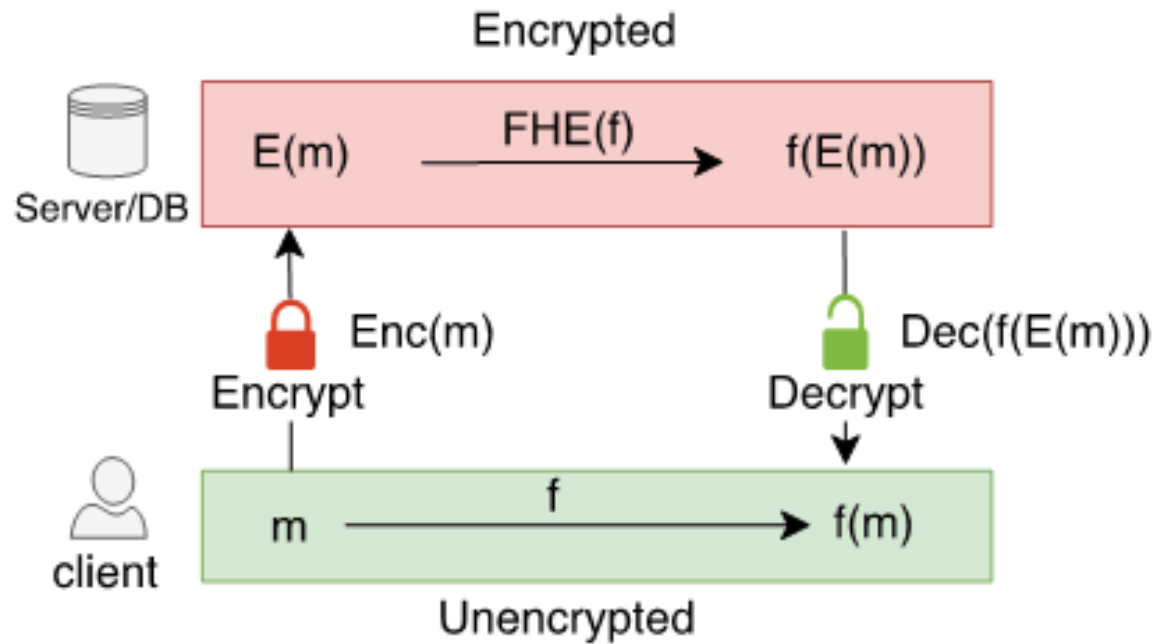
# Background - The Privacy Threats

- Online Analytical Processing (OLAP) faces privacy threats
- External attacks:
  - 2022 Uber data breach([What Caused the Uber Data Breach in 2022](#) )
  - 2024 AT&T data breach([AT&T Addresses Recent Data Set Released on the Dark Web](#))
- Internal leaks:
  - 2023 Tesla former employee information leak([Tesla insider breach exposes thousands of employees](#))
  - 2024 Evolve Bank information leak([Evolve Bank says ransomware gang stole personal data on millions of customers](#) )

# Background - Deficiencies in existing work

- The existing methods cannot meet the requirements of OLAP, or they have deficiencies in terms of efficiency or security.
  - Traditional encryption methods (e.g., AES, RSA) do not support computation.
  - Solutions based on CryptDB support limited computations and may leak data access patterns.
  - Solutions based on TEEs rely on specific hardware and are vulnerable to side-channel attacks.
  - The performance of existing methods based on Fully Homomorphic Encryption (FHE) is not good.
- Among the mentioned methods, we consider the approach based on FHE to achieve the highest level of security. Therefore, we are attempting to enhance the performance of this method.

# The Fully Homomorphic Encryption (FHE)

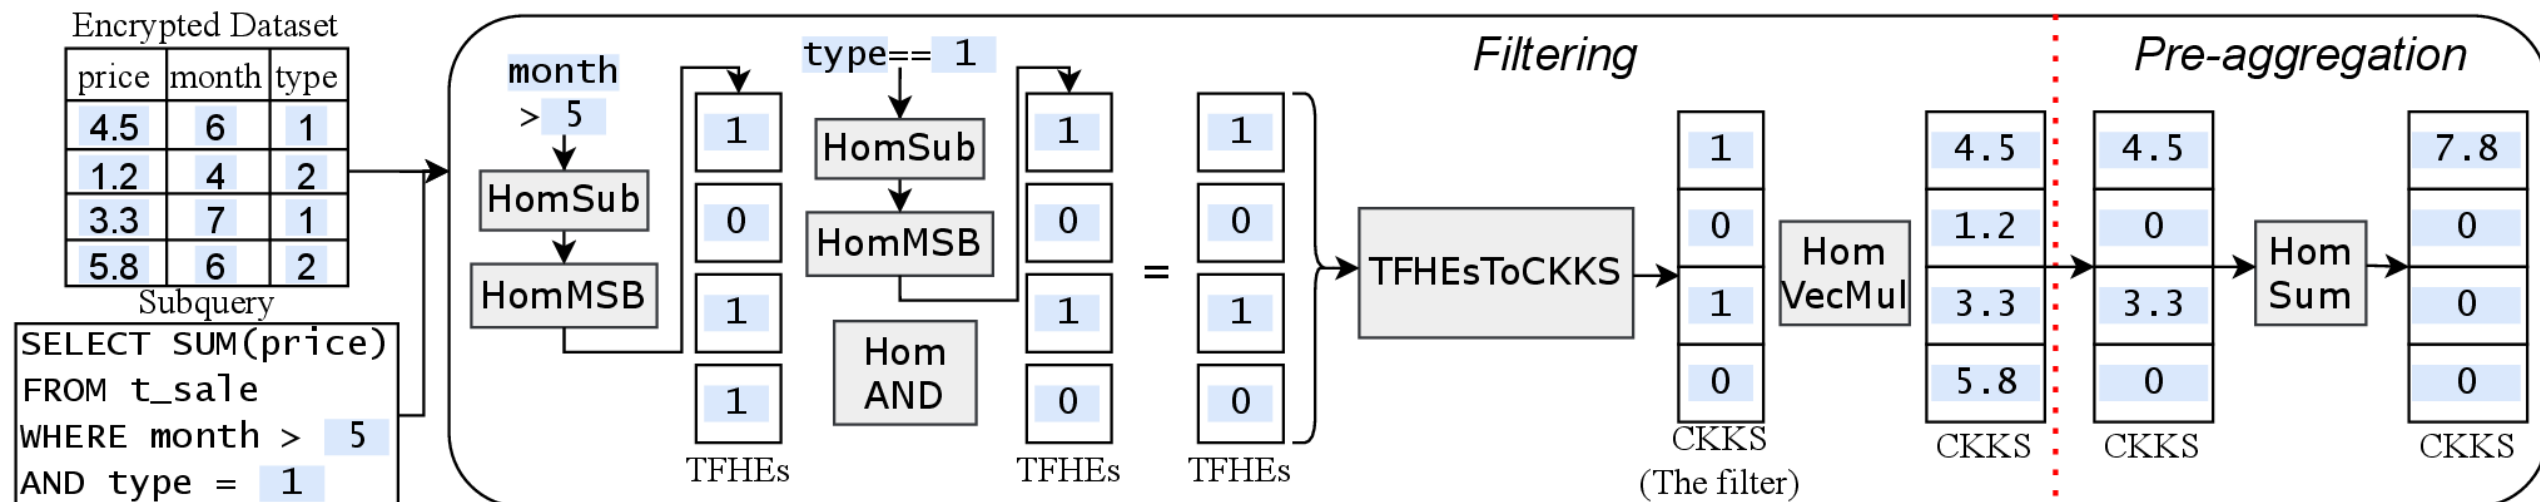


# The state-of-the-art single-machine FHE-based OLAP System – HE<sup>3</sup>DB

- HE<sup>3</sup>DB: An Efficient and Elastic Encrypted Database Via Arithmetic-And-Logic Fully Homomorphic Encryption (<https://dl.acm.org/doi/abs/10.1145/3576915.3616608>)
- Secure enough
- But with poor performance...

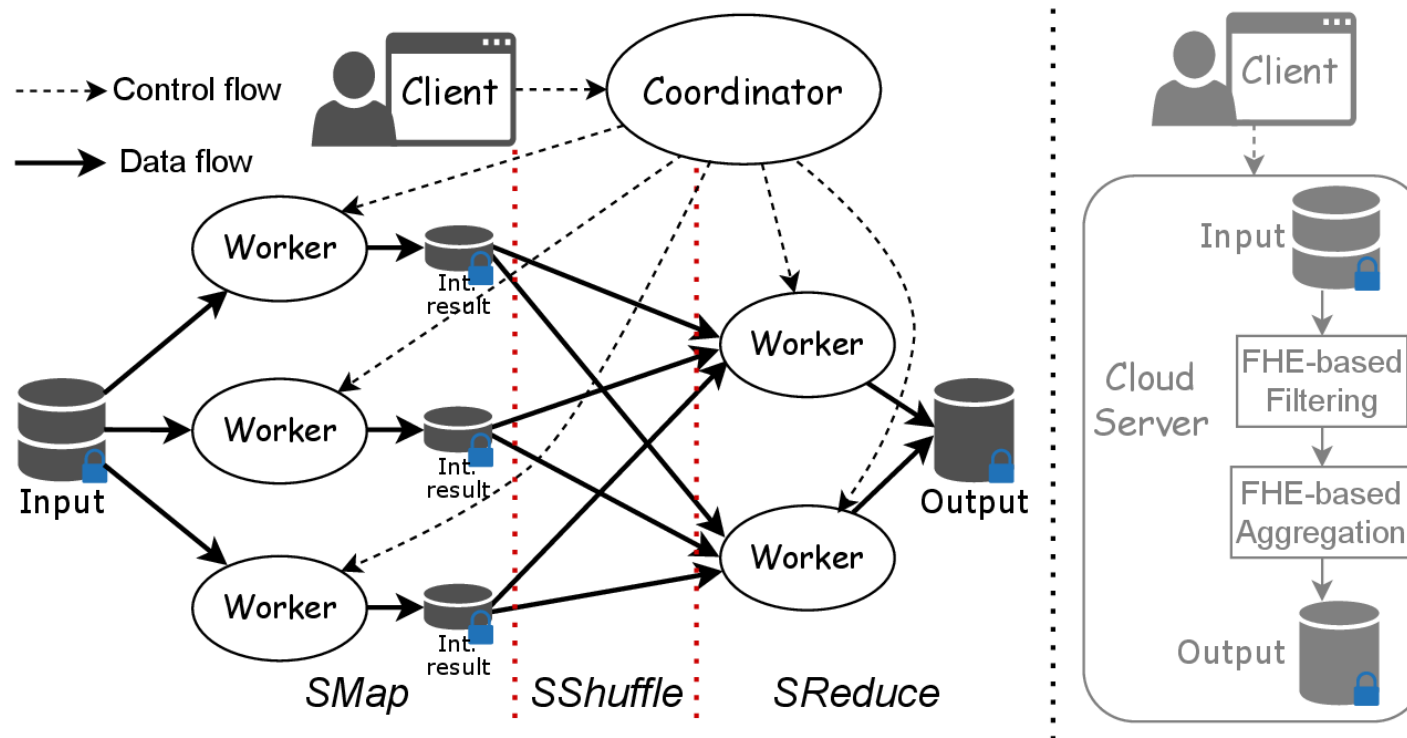
# Reasons for Slowness: The FHE-based Filtering

- FHE-based filtering:
  - Utilizing TFHE for homomorphic subtraction on all data entries
  - Performing homomorphic most significant bit (MSB) extraction
  - Conducting homomorphic bitwise AND operations
  - Transforming the results into CKKS format to generate an encrypted 0/1 vector
  - Performing homomorphic vector multiplication between the filter and the data
- Performance bottleneck: 99.96% of the overall time



# Key idea

- Can we try using parallel or distributed methods to enhance performance?
- The first step: Employing the MapReduce paradigm for distributed FHE-based OLAP



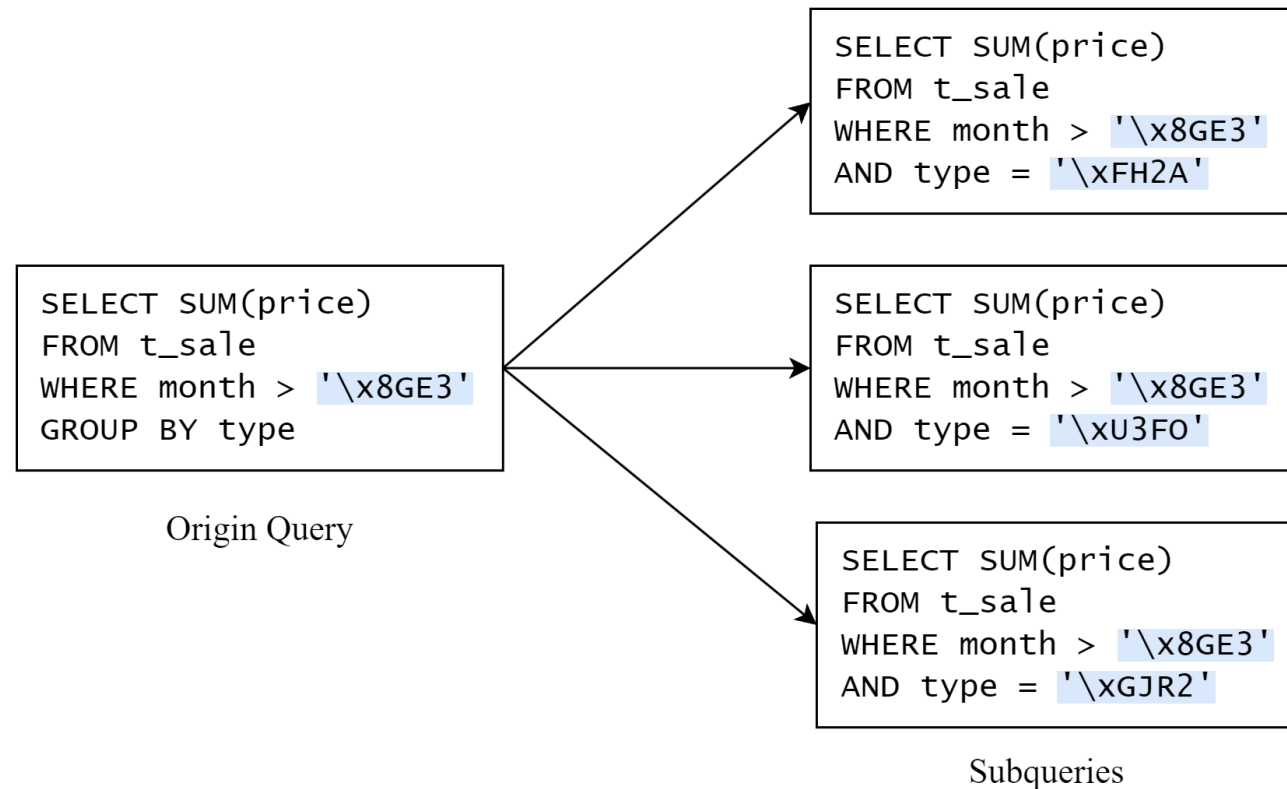
# Challenge

- The result of FHE comparison is encrypted randomly, making it impossible to group the intermediate results
- Cannot directly use MapReduce...
- Can we modify the MapReduce process based on the characteristics of FHE while ensuring security?

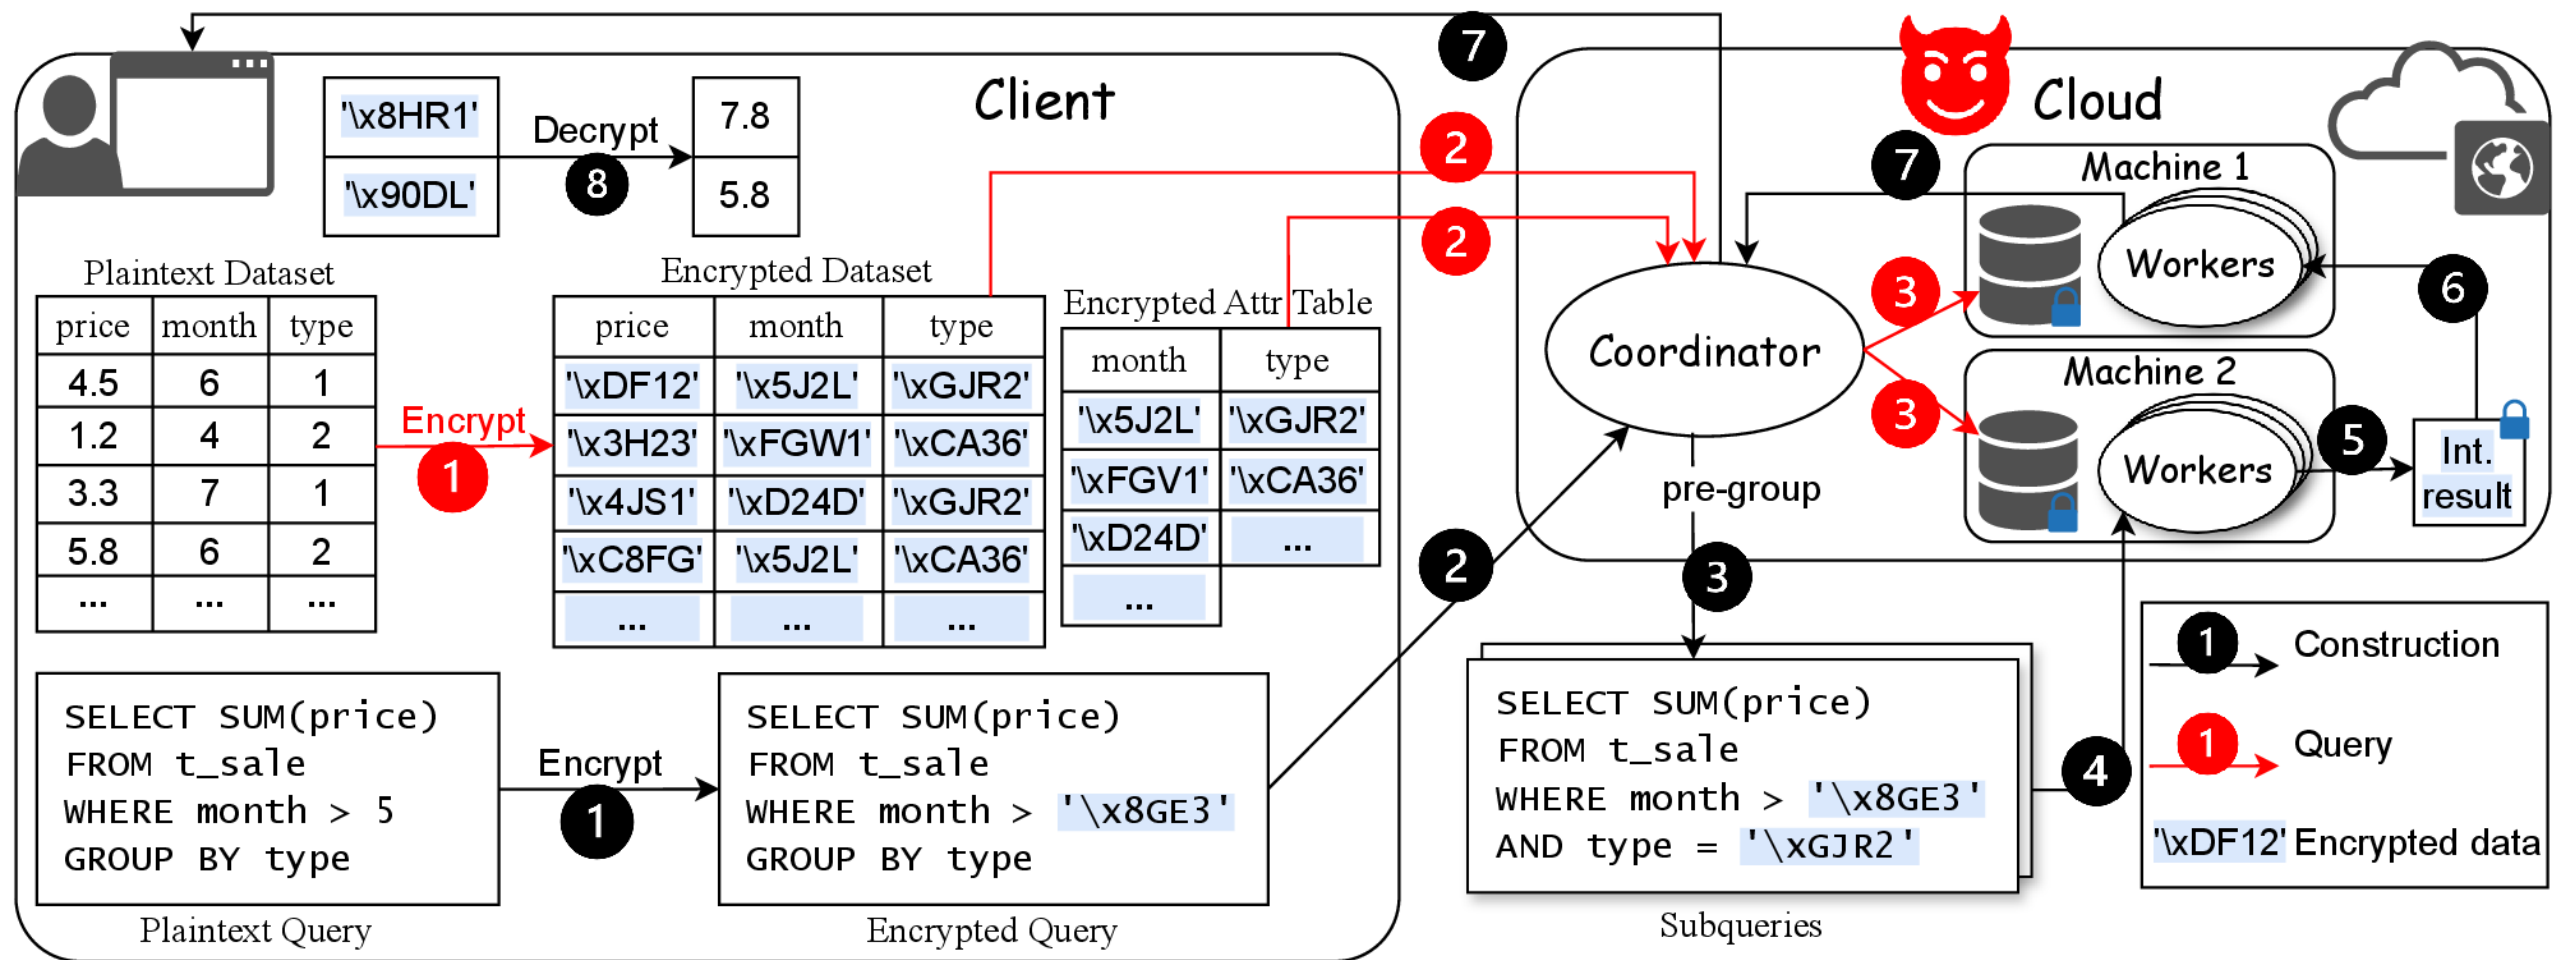


# The Pre-Group Operation

- Essentially, it involves partitioning queries in advance using attribute tables, breaking down a single query into multiple subqueries.



# The Overview of HEDAS



# System Construction

---

## Algorithm 1: System Construction

---

**Input:** Original plaintext dataset  $D_{pl}$

// Step ①: Encrypt dataset and generate attributes table

1  $pk \leftarrow \text{GetEncryptionKey}(); ek \leftarrow \text{GenerateEvalKey}(pk);$

2  $D_{en} \leftarrow \text{Encrypt}(D_{pl}, pk); attr\_table \leftarrow \text{GenerateAttributeTable}(D_{pl}, pk);$

// Step ②: Send encrypted data and attributes table to coordinator

3  $\text{SendInitData}(D_{en}, attr\_table, ek);$

// Step ③: Shard encrypted data and distribute to cloud machines

4  $machine\_num \leftarrow \text{GetMachineNum}();$

5  $shards\_array[] \leftarrow \text{DataSharding}(D_{en}, machine\_num);$

6 **For each cloud machine  $m_i$ ; do**

7   |  $\text{SendShard}(shards\_array[m_i]);$

---

# Query Processing

---

## Algorithm 2: Query Processing

---

**Input:** Plain text query scheme  $QS$  with plaintext filtering predicate parameters  $P_{pl}$  and GROUP BY attribute  $G$

```
// Step ❶: Encrypt query parameters
1  $pk \leftarrow \text{GetEncryptionKey}(); P_{en} \leftarrow \text{Encrypt}(P_{pl}, pk);$ 
// Step ❷: Send query with encrypted parameters to coordinator
2  $\text{SendQuery}(QS, P_{en}, G);$ 
// Step ❸: pre-group operation
3  $\text{attr\_table} \leftarrow \text{GetAttributeTable}();$ 
4  $\text{sub\_queries}[] \leftarrow \text{PreGroup}(QS, P_{en}, G, \text{attr\_table});$ 
// Step ❹: Assign SMap tasks to worker nodes
5 While SMap not finished; do
6    $\text{worker} \leftarrow \text{WaitWorker}(); SMap\_task \leftarrow \text{GetSMapTask}(\text{sub\_queries});$ 
7    $\text{AssignTask}(\text{worker}, SMap\_task);$ 
// Step ❺: Workers process SMap tasks
8  $SMap\_task \leftarrow \text{RequestSMapTask}();$ 
9  $\text{int\_result} \leftarrow \text{DoSMapTask}(SMap\_task); \text{CacheIntResult}(\text{int\_result});$ 
// Step ❻: Assign SReduce tasks
10 While SReduce not finished; do
11    $\text{worker} \leftarrow \text{WaitWorker}(); SReduce\_task \leftarrow \text{GetSReduceTask}();$ 
12    $\text{AssignTask}(\text{worker}, SReduce\_task);$ 
// Step ❼: Workers process SReduce tasks
13  $SReduce\_task \leftarrow \text{RequestSReduceTask}();$ 
14  $\text{int\_results} \leftarrow \text{GetIntResults}(SReduce\_task.key);$ 
15  $\text{partial\_result} \leftarrow \text{DoSReduceTask}(SReduce\_task, \text{int\_results});$ 
16  $\text{SendResultToCoordinator}(\text{partial\_result});$ 
// Step ❽: Coordinator collects and returns the final result
17  $\text{result} \leftarrow \text{CollectAllResults}(); \text{ReturnResult}(\text{result});$ 
```

---

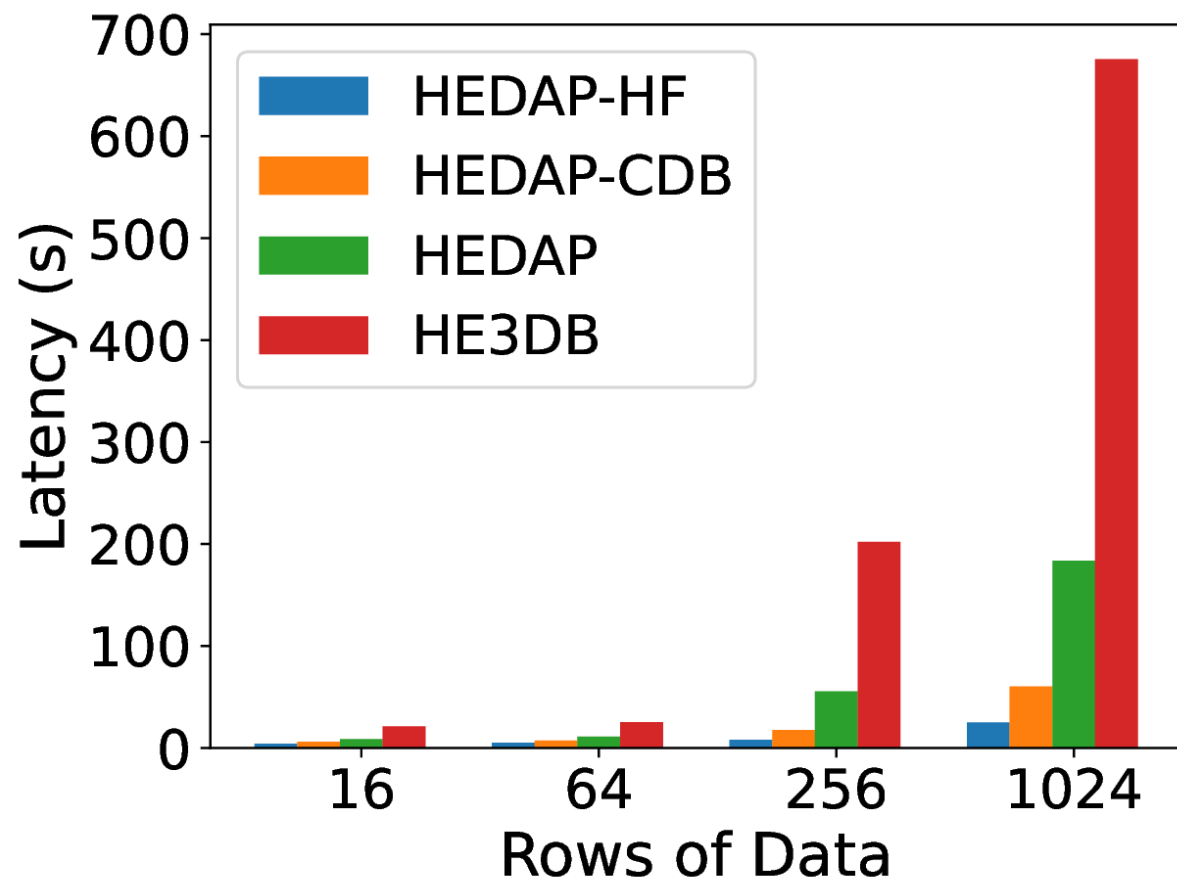
# The case studies

- Our system can integrate traditional CryptDB indexes or a type of tree-like HashFilter-based indexes (another work of ours) in scenarios where privacy requirements are not as stringent.
- Two case study subsystems: HEDAS-CryptDB and HEDAS-HashFilter.

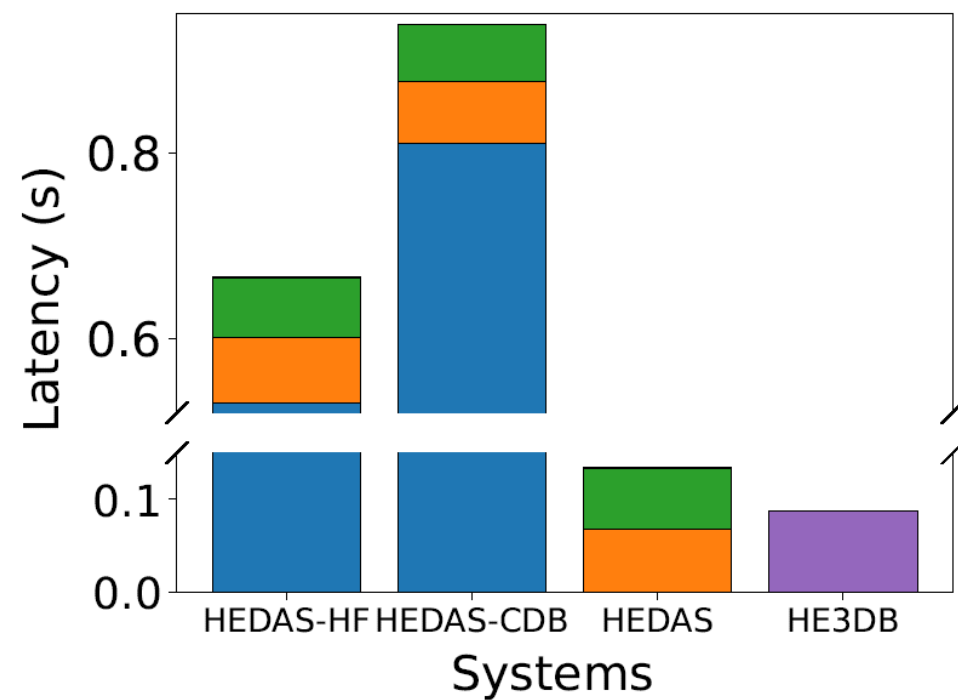
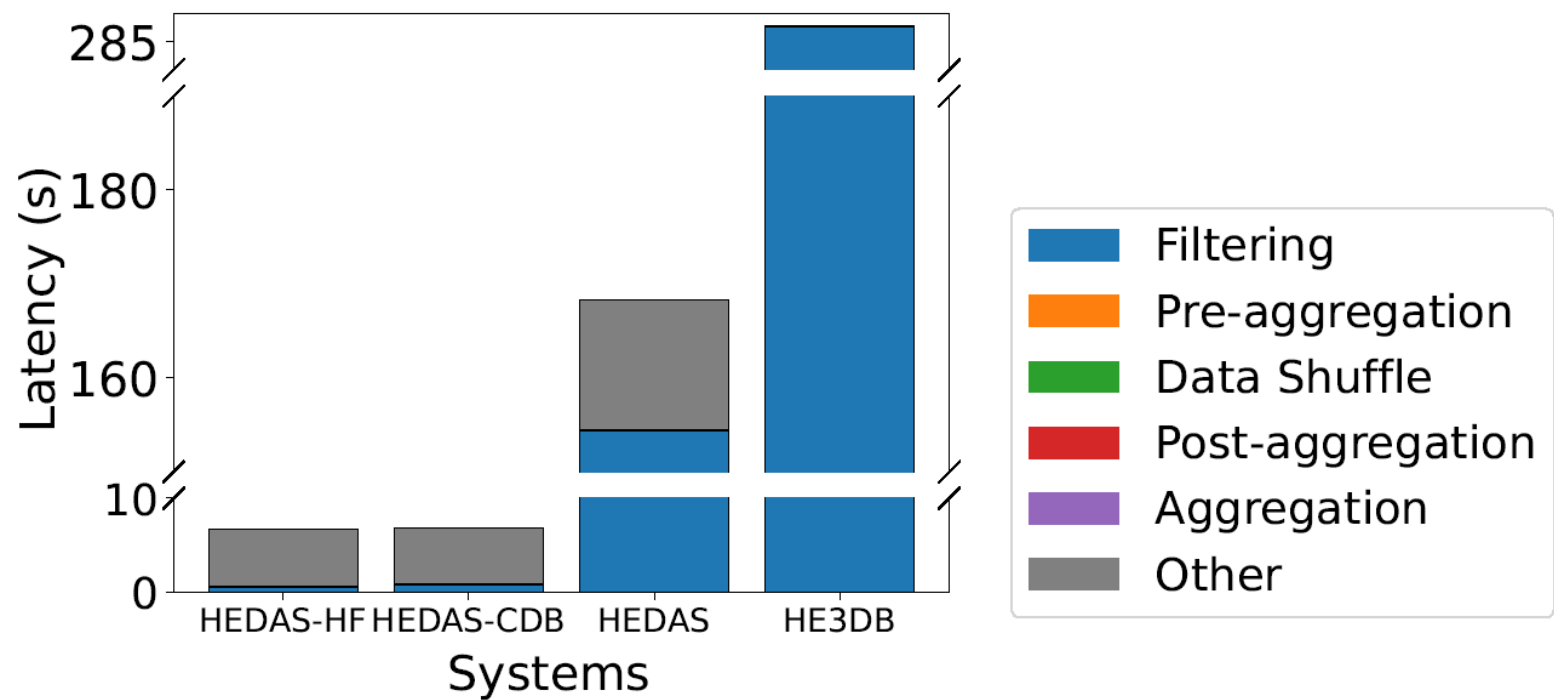
# Evaluation questions

- What is the end-to-end and breakdown performance gain compared to HE3DB?
- Additional performance improvement for the case study systems?
- How about the scalability?

# The End-to-end Latency

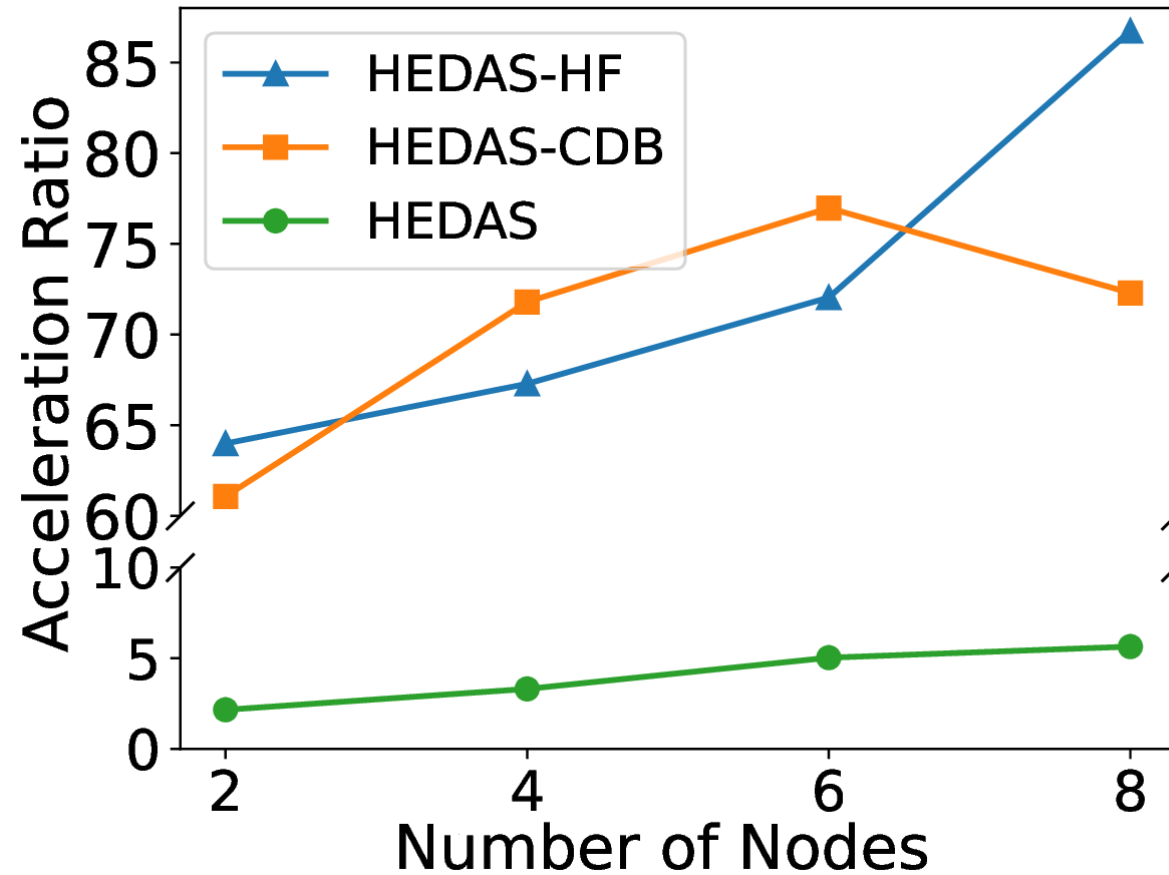


# The Breakdown latency





# The Scalability with Increasing Nodes



# Conclusion & Future work

- HEDAS is the first distributed FHE-based OLAP system in the MapReduce style, enabling secure and efficient OLAP operations.
- We will continue to explore systematic methods to optimize FHE-based databases, and HEDAS is just the beginning.
- Our future work includes:
  - Utilizing accelerators (e.g., GPUs) to provide more powerful parallel computing;
  - Designing special cache mechanisms to accelerate FHE filtering;
  - Developing more efficient indexing structures for FHE-based databases.

**THANK YOU!**