

# TSM-Bench: Benchmarking Time Series Database Systems for Monitoring Applications

Abdelouahab Khelifati, Mourad Khayati, Anton Dignös,  
Djellel Difallah, and Philippe Cudré-Mauroux

## Goal and Contribution

**Motivation:** Existing Time Series Database Systems (TSDBs) benchmarks are limited in the number of evaluated systems, the type of workloads, the size and type of data, and the query variability.

**Goal:** A comprehensive benchmark of TSDBs for monitoring applications.

**Contributions:**

1. Extensive evaluation of seven popular TSDBs using temporal workloads.
2. A new time series generation technique.
3. Recommendations for understanding and navigating systems' architectural designs.

## Applications

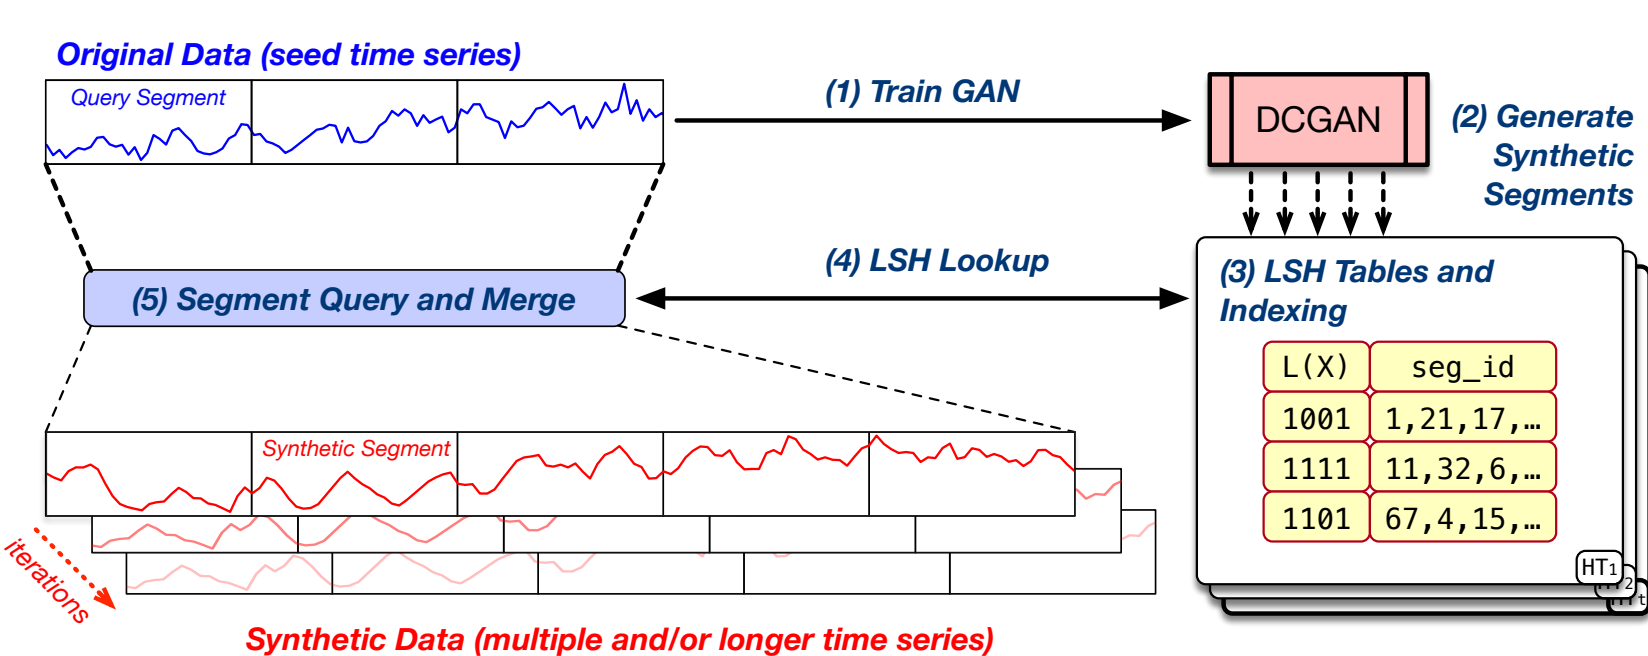
**Monitoring of Watercourse (BAFU)**

- BAFU monitors the water discharge and level in Swiss rivers.
- Evaluates water quality.
- Assesses the impact of climate change and triggers alerts in case of hazard.

**Other applications:** Internet of Things (IoT), smart grids, traffic networks, etc.

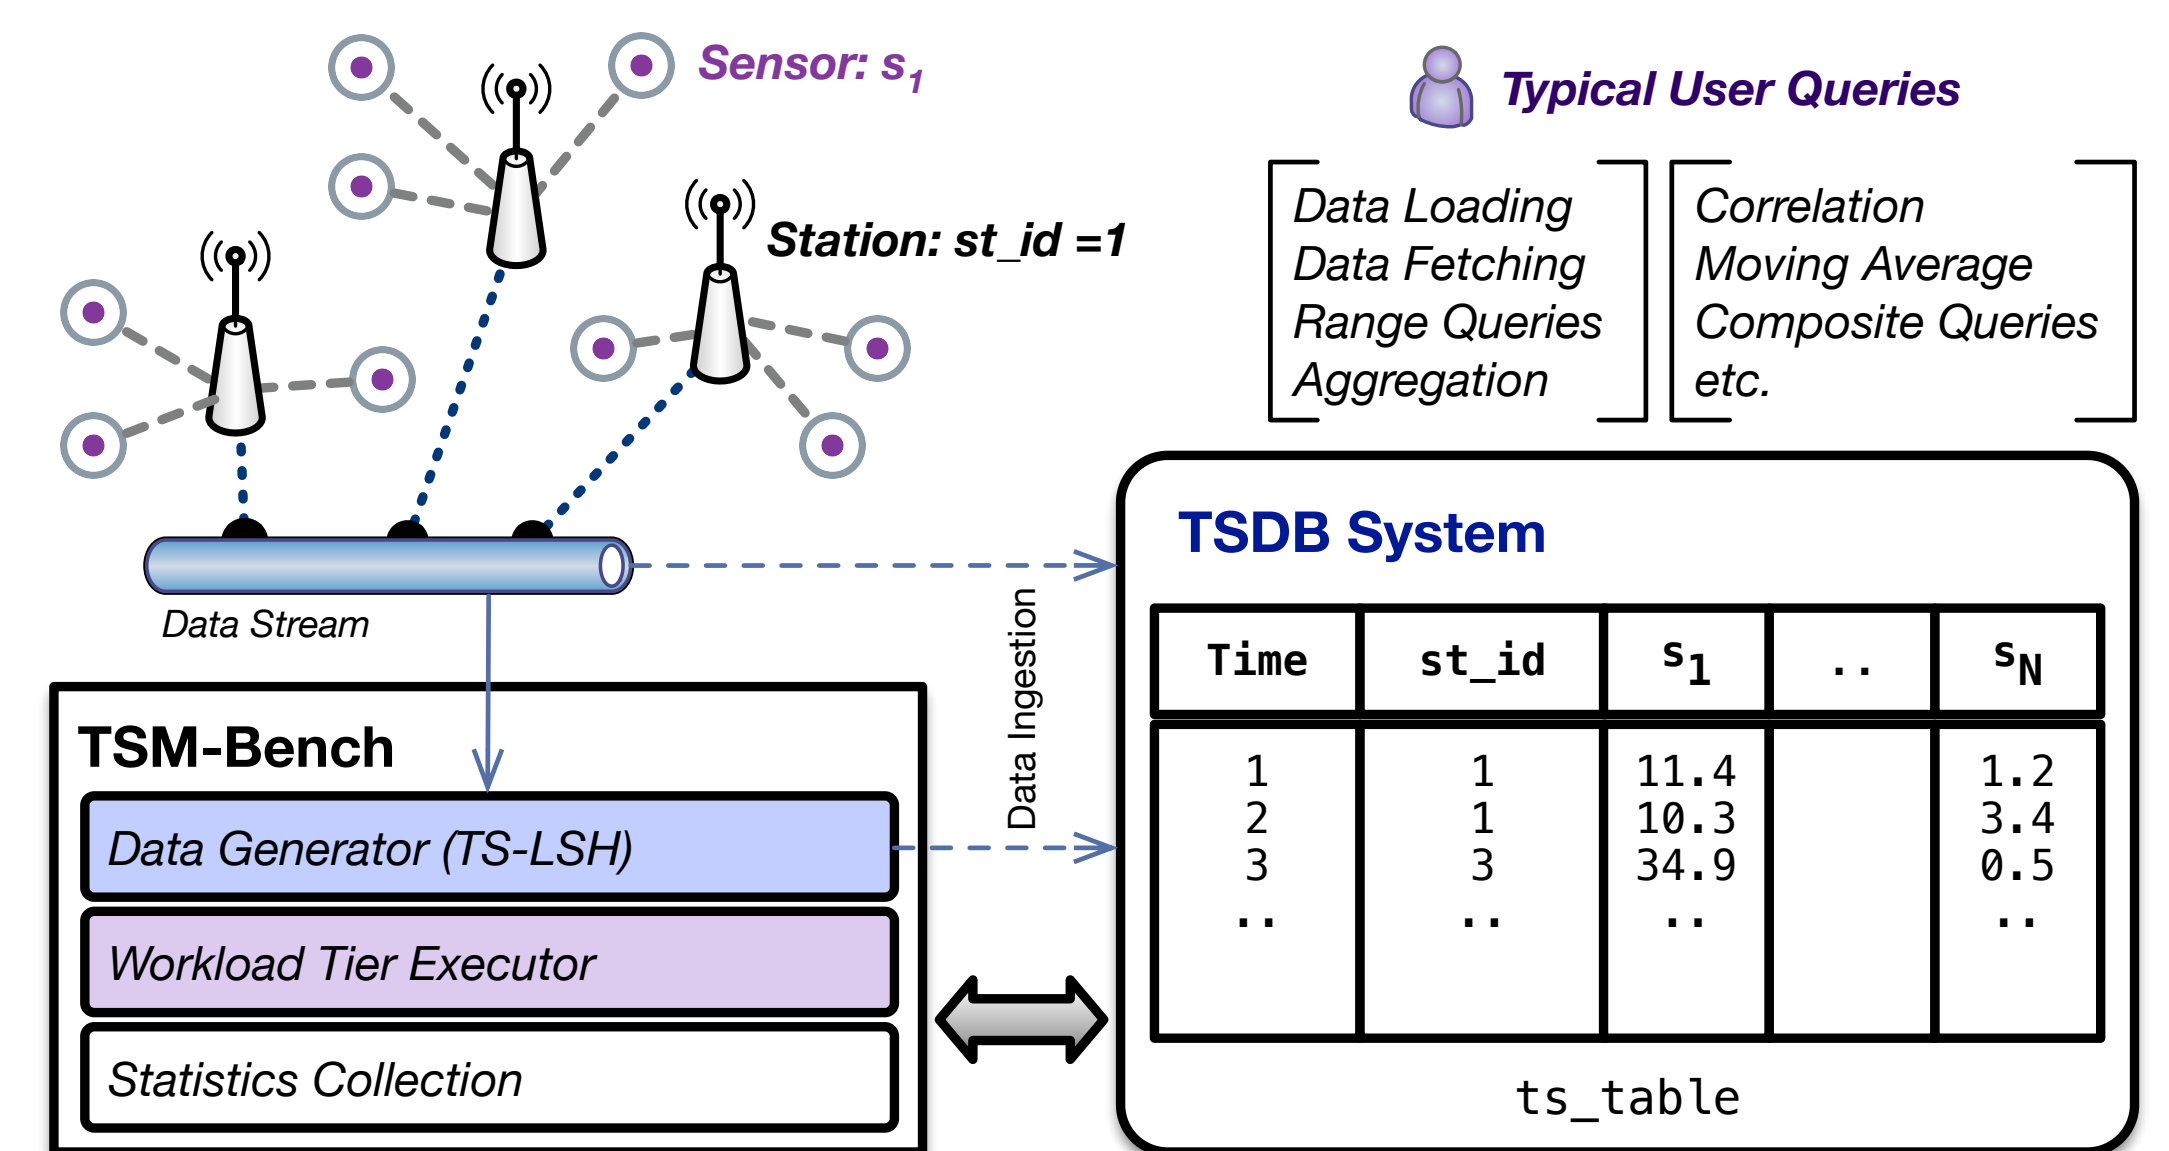
## Time Series Generator

- A new generation technique that combines GAN with LSH.
- Scalable data generation of large realistic time series.




## TSM Architecture

- TS-LSH uses sample data to generate large data streams.
- The executor launches configurable workload tiers.
- The statistics collection module records the performance of the TSDB.



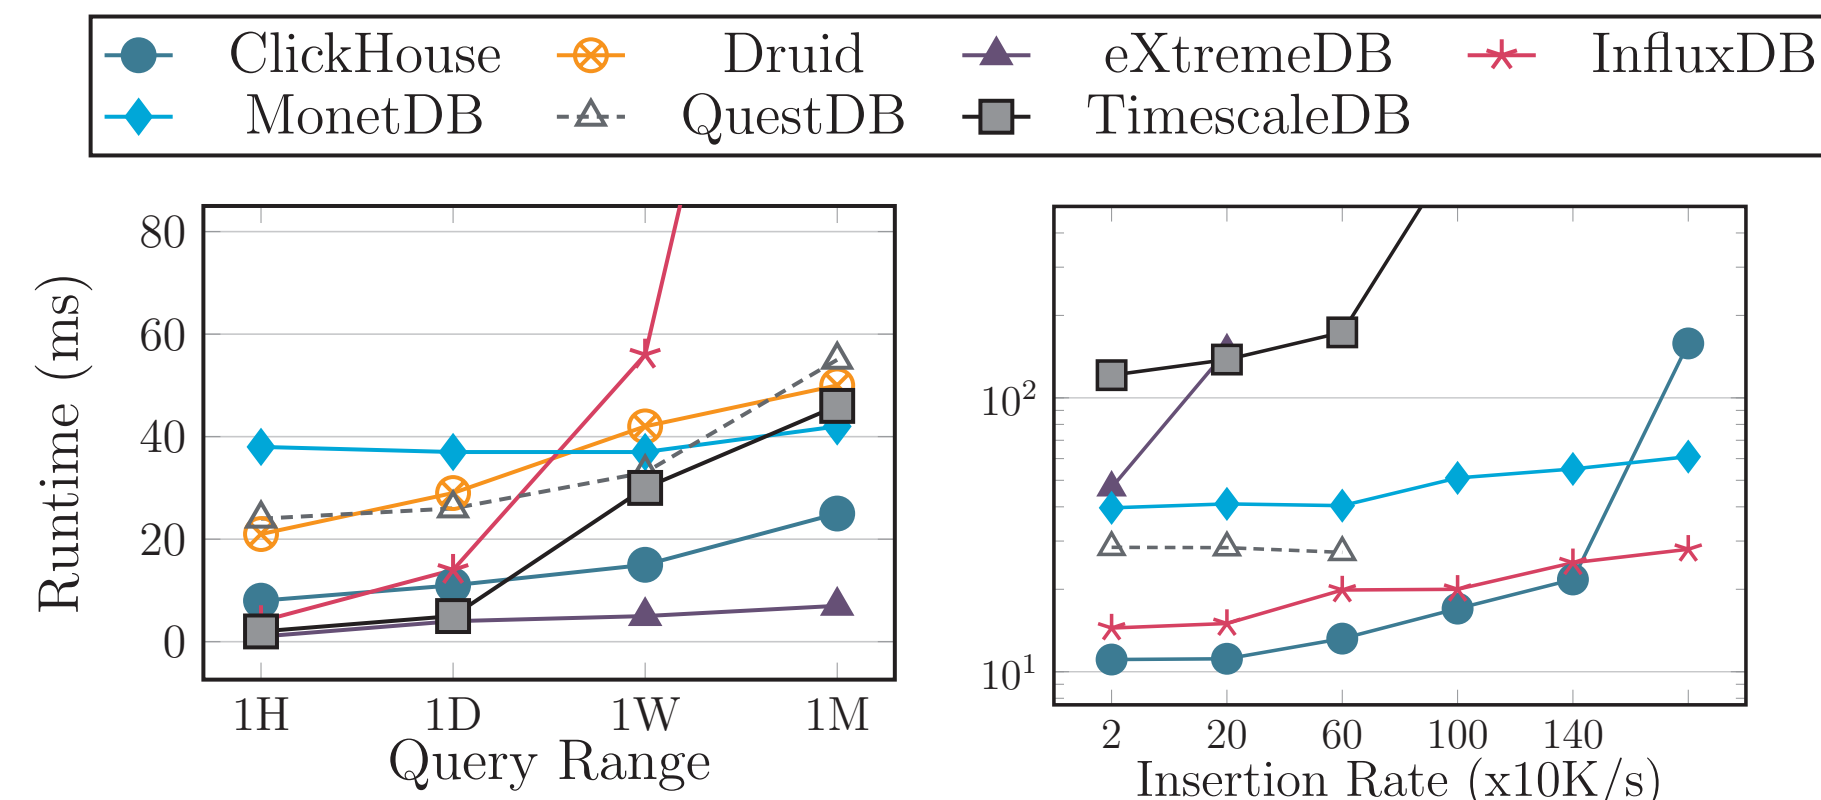
## Experiments

- The performance of the systems depends on the size of the input/output data.
- The offline and online workloads show different trends.
- Time series features heavily impact systems' compression capability.



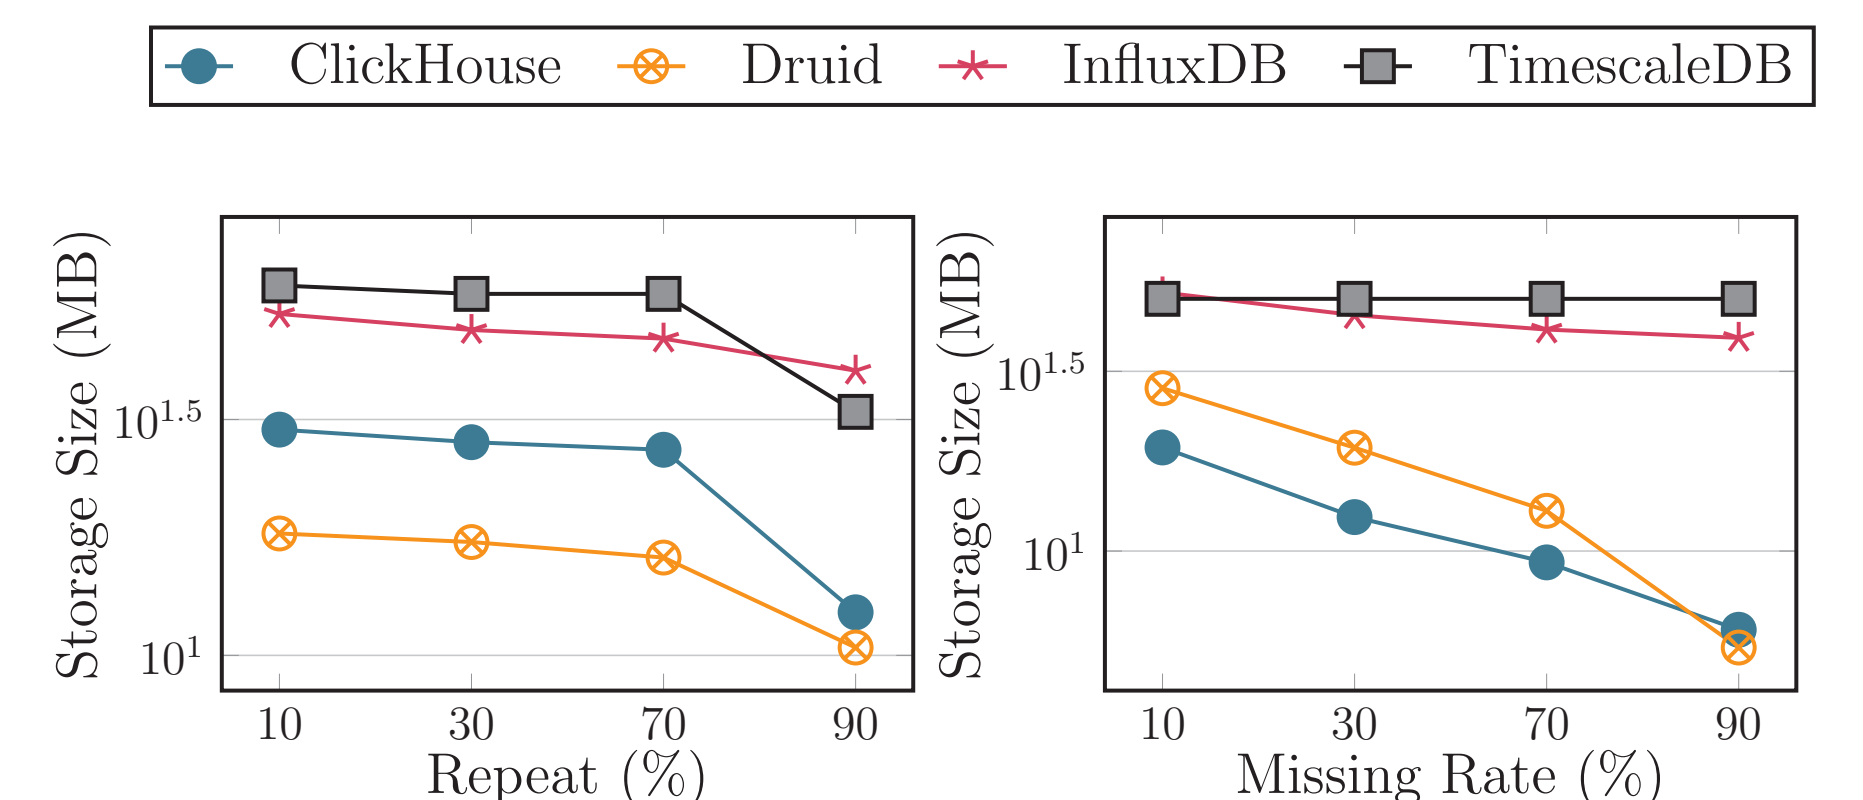
(a) Fetching.

(b) Aggregation.



(a) Offline mode.

(b) Online mode.

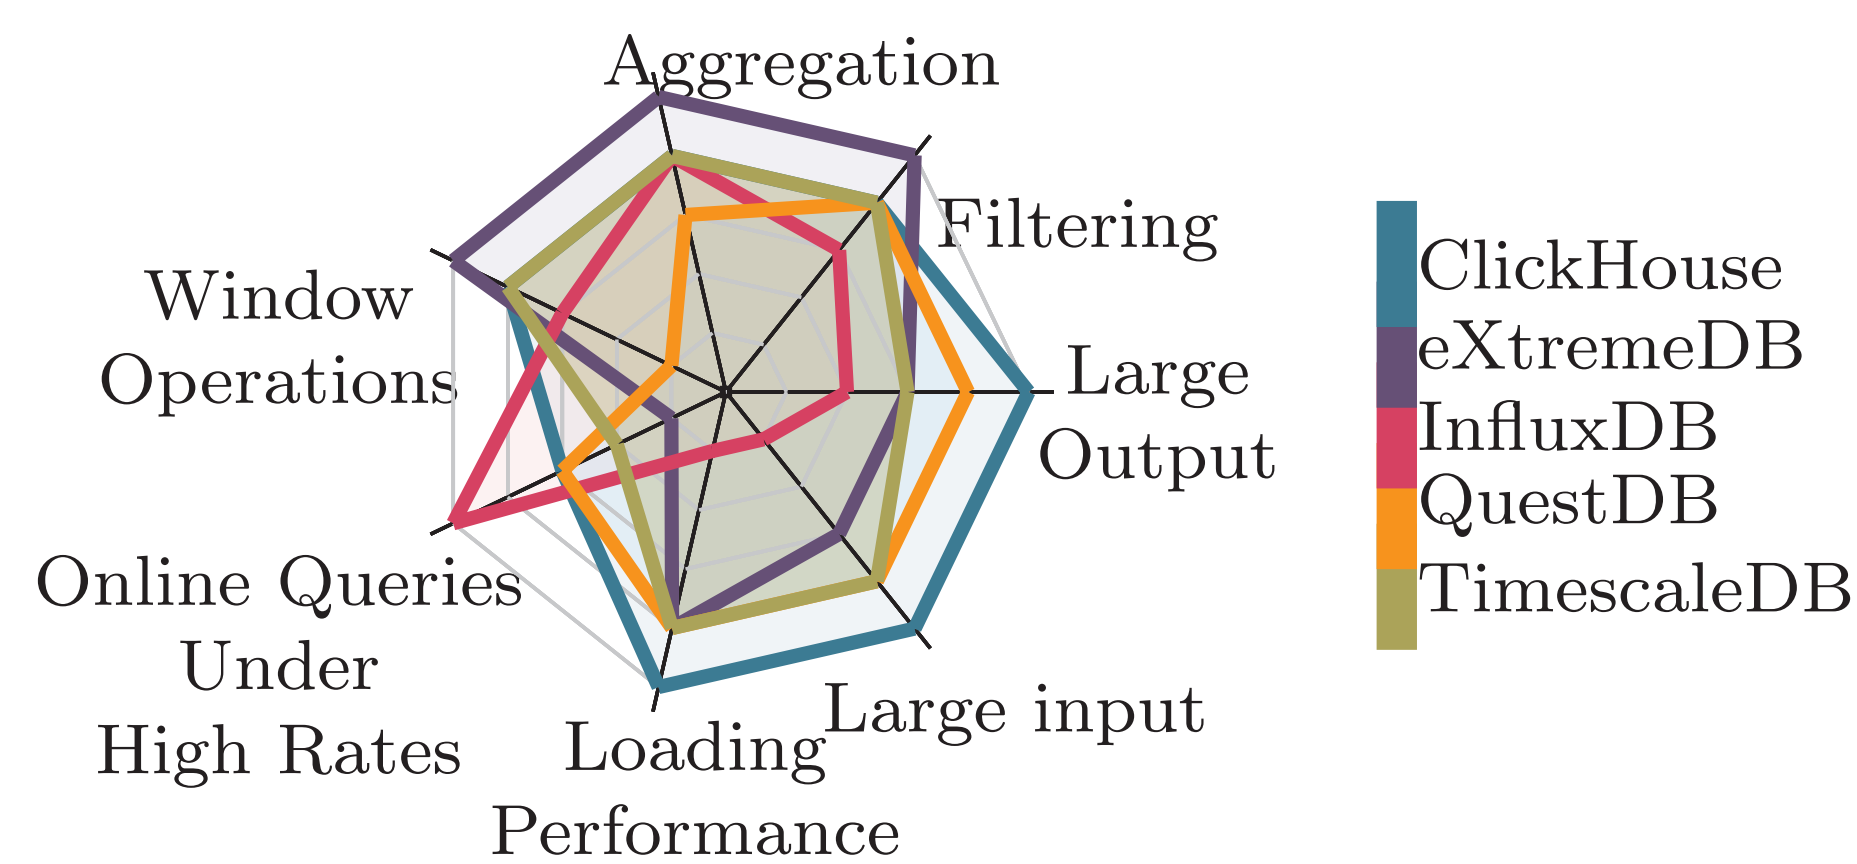


(a) Repeats.

(b) Data Sparsity.

## Performance Summary

- Seven discriminative dimensions.
- Performance ranking for different query types on a 0-5 scale.
- No silver bullet.
- Clickhouse and extremeDB offer the best trade-off.



## Additional Info

- Github: <https://github.com/eXascaleInfoLab/TSM-Bench>
- Related works:
  - Difallah D., Pavlo A., Curino C., and Cudré-Mauroux P.: "OLTP-Bench: An Extensible Testbed for Benchmarking Relational Databases", VLDB 2013.
  - Khayati M., Lerner A., Tymchenko Z., and Cudré-Mauroux P.: "Mind the Gap: An Experimental Evaluation of Imputation of Missing Values Techniques in Time Series", PVLDB 2020.

## Design Choices

