



PDF Download  
3722212.3724446.pdf  
06 January 2026  
Total Citations: 0  
Total Downloads: 765

Latest updates: <https://dl.acm.org/doi/10.1145/3722212.3724446>

RESEARCH-ARTICLE

## OpenMLDB: A Real-Time Relational Data Feature Computation System for Online ML

XUANHE ZHOU, Shanghai Jiao Tong University, Shanghai, China

WEI ZHOU, Shanghai Jiao Tong University, Shanghai, China

LIGUO QI

HAO ZHANG

DIHAO CHEN

BINGSHENG HE, National University of Singapore, Singapore City, Singapore

[View all](#)

Open Access Support provided by:

[Shanghai Jiao Tong University](#)

[Tsinghua University](#)

[National University of Singapore](#)

Published: 22 June 2025

[Citation in BibTeX format](#)

SIGMOD/PODS '25: International  
Conference on Management of Data  
June 22 - 27, 2025  
Berlin, Germany

Conference Sponsors:  
SIGMOD

# OpenMLDB: A Real-Time Relational Data Feature Computation System for Online ML

Xuanhe Zhou  
Shanghai Jiao Tong  
University  
Shanghai, China  
zhouxh@cs.sjtu.edu.cn

Wei Zhou  
Shanghai Jiao Tong  
University  
Shanghai, China  
weizhoudb@sjtu.edu.cn

Liguo Qi  
4Paradigm  
Shanghai, China  
qiliguo@4paradigm.com

Hao Zhang  
4Paradigm  
Beijing, China  
zhanghao@4paradigm.com

Dihao Chen  
SF Express  
Shenzhen, China  
chendh@sf-express.com

Bingsheng He  
National University of  
Singapore  
Singapore, Singapore  
hebs@comp.nus.edu.sg

Mian Lu  
4Paradigm  
Singapore, Singapore  
lumian@4paradigm.com

Guoliang Li  
Tsinghua University  
Beijing, China  
liguoliang@tsinghua.edu.cn

Fan Wu  
Shanghai Jiao Tong  
University  
Shanghai, China  
fwu@cs.sjtu.edu.cn

Yuqiang Chen  
4Paradigm  
Beijing, China  
chenyuqiang@4paradigm.com

## Abstract

Efficient and consistent feature computation is crucial for a wide range of online ML applications. Typically, feature computation is divided into two distinct phases, i.e., offline stage for model training and online stage for model serving. These phases often rely on execution engines with different interface languages and function implementations, causing significant inconsistencies. Moreover, many online ML features involve complex time-series computations (e.g., functions over varied-length table windows) that differ from standard streaming and analytical queries. Existing data processing systems (e.g., Spark, Flink, DuckDB) often incur multi-second latencies for these computations, making them unsuitable for real-time online ML applications that demand timely feature updates.

This paper presents OpenMLDB, a feature computation system deployed in 4Paradigm’s SageOne platform and over 100 real scenarios. Technically, OpenMLDB first employs a *unified query plan generator* for consistent computation results across the offline and online stages, significantly reducing feature deployment overhead. Second, OpenMLDB provides an *online execution engine* that resolves performance bottlenecks caused by long window computations (via pre-aggregation) and multi-table window unions (via data self-adjusting). It also provides a high-performance *offline execution engine* with window parallel optimization and time-aware data skew resolving. Third, OpenMLDB features a *compact data format*

and *stream-focused indexing* to maximize memory usage and accelerate data access. Evaluations in testing and real workloads reveal significant performance improvements and resource savings compared to the baseline systems. The open community of OpenMLDB now has over 150 contributors and gained 1.6k stars on GitHub<sup>1</sup>.

## CCS Concepts

• Information systems → Relational database model.

## Keywords

Relational Data Feature Computation, Unified Query Plan Generator, Compact Time-series Data Format

### ACM Reference Format:

Xuanhe Zhou, Wei Zhou, Liguo Qi, Hao Zhang, Dihao Chen, Bingsheng He, Mian Lu, Guoliang Li, Fan Wu, and Yuqiang Chen. 2025. OpenMLDB: A Real-Time Relational Data Feature Computation System for Online ML. In *Companion of the 2025 International Conference on Management of Data (SIGMOD-Companion ’25)*, June 22–27, 2025, Berlin, Germany. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3722212.3724446>

## 1 Introduction

Relational data feature computation, which extracts ML features from raw data tables, is essential for many online applications leveraging machine learning models (e.g., credit card anti-fraud, online advertising [28, 55, 59]). However, this task is inherently complex. For instance, the product recommendation model at Vipshop requires over 600 distinct features [35]. Many of these features demand specialized ML functions, such as statistical functions (e.g., top-N frequency counts) and feature signatures (e.g., export format labeling for high-dimension features like product items [46]), as well as computation patterns (e.g., multi-table windows) distinct from those in standard streaming and analytical queries [64, 66].

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

*SIGMOD-Companion ’25*, June 22–27, 2025, Berlin, Germany

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1564-8/2025/06.

<https://doi.org/10.1145/3722212.3724446>

<sup>1</sup><http://github.com/4paradigm/OpenMLDB>

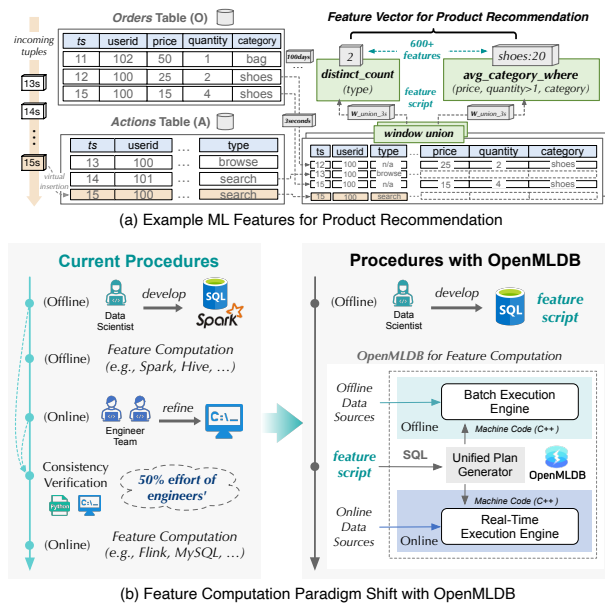


Figure 1: Relational Data Feature Computation.

**Example.** To illustrate the complexity, consider the product recommendation scenario that generates personalized advertisements based on the user’s most recent actions. As shown in Figure 1 (a), the feature computation process begins by computing window tables. For instance, we retrieve tuples from Actions and Orders that match the incoming tuple’s user ID and fall within a short time interval (e.g., three seconds). These matched tuples, together with the incoming tuple, form a window table ( $w\_union\_3s$ ). Next, we compute various window features like (i) counting the number of product types purchased by the user and (ii) calculating the average price of products within specific categories, subject to conditions like “purchased more than one”. To support recommendation at scale, feature computation must handle extremely high concurrency (e.g., more than 200 million requests per minute) and still deliver up-to-date features within milliseconds.

In current practice, a feature computation workflow has two separate stages. As shown in Figure 1 (b), first, data scientists develop and compute offline features with batch processing engines (e.g., SparkSQL), iterating until the model’s training performance meets their objectives. Afterward, to deploy these offline feature scripts in online production, the engineering team must refactor and optimize them for online computation (with streaming engines like Flink). However, the two stages are handled by different teams and systems, which can cause inconsistencies. For instance, in a real-world de-fraud case at Varo [11], the online application defines the “account balance” feature as the current real-time balance, while offline data scientists define it as the previous day’s balance. Such discrepancies cause mismatches between offline and online results, making extensive and costly verification necessary before the final deployment. Based on our experience (e.g., at Akulaku [14]), consistency verification can take several months or even one year.

Moreover, existing data processing systems, both online [1, 9, 10, 23–25, 36, 52, 54] and offline [2, 8, 39, 53, 58, 62], as well as specialized feature stores [4–6, 27], struggle to meet the performance demands (see Section 2.2). For instance, streaming execution

engines like Apache Flink [36], and Spark Streaming [52] lack critical optimizations for feature computation, such as handling long window operations and aggregating data across multiple table windows, making it difficult to achieve sub-second online response time. Similarly, widely used storage engines (e.g., MySQL [24], PostgreSQL [25], Cassandra [1]) are not designed for rapid ingestion or efficient ordering of time-series data, both of which are essential for online feature storage. On the offline side, the execution engines (e.g., Spark [62], Hive [58], ClickHouse [53]) face issues including limited operation-level parallelism, data skew, and frequent RPC calls during distributed computation, which limit their throughput and scalability for offline feature computation. Lastly, existing feature stores like Feast [4] primarily focus on fast data retrieval rather than real-time feature computation, forcing users to combine multiple tools for feature computation and thus introducing additional synchronization and operational overhead.

To address these problems, we propose OpenMLDB, an industrial-level feature computation system. Specifically, OpenMLDB first ensures **online-offline feature computation consistency** by providing a unified query plan generator that supports extended SQL syntax and compilation-level optimizations (e.g., cycle binding, compilation caching). Second, OpenMLDB achieves **ultra-high feature computation performance** through a suite of optimization techniques for both online and offline computations. For online computation, it reuses pre-aggregation results to enhance long-window computations and employs a self-adjusting strategy to optimize multi-table window unions. For offline computation, it supports multi-window parallel optimization and mitigates data skew by repartitioning tuples based on utilized columns and data distribution. Finally, OpenMLDB offers a compact in-memory data format and a stream-focused data structure, which pre-ranks data by key and timestamp to enable rapid online data access.

We summarize our main contributions as follows.

- We present OpenMLDB, an industry-level feature computation system recognized as a recommended feature product by Gartner [41], which has been deployed in 100+ real-world scenarios<sup>2</sup>.
- We provide a unified query plan generator that streamlines feature script writing for offline and online stages and reduces feature deployment time from months to days or even less.
- We propose an online real-time execution engine that resolves performance bottlenecks in ML feature computations such as long window aggregations and multi-table window unions.
- We propose an offline batch execution engine that tackles challenges such as low parallelism in multi-window aggregations and data-skewed partitionings in existing systems.
- We provide a compact in-memory data format for efficient memory usage and stream-oriented structures to optimize frequent data access during online feature computation.
- We thoroughly evaluated OpenMLDB across various workloads. Our results show that OpenMLDB delivers (1) 10x–20x higher online performance than Flink and DuckDB; (2) 6x faster offline performance than Spark and MPP databases like GreenPlum; and (3) 40% lower memory usage compared to in-memory databases like Redis.

<sup>2</sup>See user feedback at <https://github.com/4paradigm/OpenMLDB/discussions/707>.

## 2 Background and Related Work

### 2.1 Relational Data Feature Computation

Given a set of data tables  $\mathcal{R}$ , the goal of relational data feature computation is to produce feature values, as specified in a feature script  $Q$ , that may include both discrete and continuous features. These feature computations often extend beyond standard relational operations (see Section 4.1) and present unique challenges. For example, they may require specialized windowed aggregations (e.g., top-frequency keys, conditional averages, time-series indicators), advanced string parsing, window unions, last-join operations, and direct output in ML-friendly formats. Based on distinct performance requirements, it can be further divided into online and offline stages.

**Online Feature Computation** processes features for continuously arriving data tuples and provides up-to-date features for online requests. The computations draw on both recent data (often spanning windows of different time lengths) and static reference data (e.g., user profiles). Many commercial scenarios demand stringent real-time performance. For example, among the target clients of OpenMLDB, *Akulaku* (a fintech company) requires about 4-ms response time for risk control, and bank’s real-time anti-fraud checks must be completed under 20 ms [16]. Similarly, companies such as *UBiX* and *Vipshop* need sub-100 ms latency for online advertising.

**Offline Feature Computation** processes features in batch mode against periodically updated tables (e.g., historical orders). The primary objective is to efficiently handle large batches of computation requests while maximizing system throughput. For instance, in consumer-facing scenarios like banks and *Yum Holdings* (a restaurant company), daily active users reach tens of millions, and achieving higher throughput helps reduce machine costs. In a similar vein, when *Akulaku* adopted OpenMLDB for offline risk control, they saved 20 servers (each with 128GB RAM), totaling around \$275,000.

### 2.2 Existing Data Systems

**Online Execution Engines.** General-purpose stream processing systems [36, 52] are insufficient for millisecond-level online ML services. Apache Flink [36] incurs significant overhead due to frequent context switching and lacks advanced optimizations for stateful aggregations. For instance, with sliding windows, Flink needs to select and discard the oldest data during each window operation. Since it lacks a robust state retention mechanism, Flink has to re-sort the data to identify the oldest entries for eviction, increasing the time complexity from  $O(1)$  to  $O(\log n)$ . Additionally, its Java-based implementation lacks specific low-level optimizations, resulting in high latency for individual requests. Similarly, Spark Streaming [52] relies on multi-node computation and requires frequent RPC calls to complete operations, making it challenging to achieve latencies in the tens-of-milliseconds range.

**Online Storage Engines.** Relational databases such as MySQL [24], MariaDB [10], and PostgreSQL [25], even when configured for in-memory storage engines, often exhibit inefficiencies in standard indexing for time-based queries (e.g., rapid data insertion). They also lack native time-ordering capabilities essential for real-time analytics and sequential data processing. Distributed storage systems like HBase [9], Cassandra [1], DynamoDB [54], and

MongoDB [23] do not support standard ANSI SQL and lack optimizations for time-series data, which increases usage complexity for data scientists. While NewSQL systems (e.g., TiDB [43], CockroachDB [56], OceanBase [61]) simplify usage via SQL support, they are also not optimized for time-series data. Meanwhile, time-series databases (e.g., InfluxDB [21], TDengine [18], OpenTSDB [17]) cannot fully manage data in memory space (primarily on disk), limiting their ability to support high-performance feature extraction.

**Offline Execution Engines.** Existing data processing systems [39, 58, 62] lack specific optimizations for feature extraction. For instance, Spark’s window aggregation does not support Whole-Stage Java Code Generation and struggles to parallelize many window operations without overlaps. Similarly, Apache Hive [58], built on top of MapReduce [39], suffers from high job startup latency and disk-based operations, rendering it unsuitable for high-throughput iterative feature extraction. Other MPP systems, including Apache Doris [2], ClickHouse [53], GreenPlum [8], Apache Druid [3], and StarRocks [26], are primarily designed for general analytical workloads rather than time-series feature computations. Furthermore, existing *C++ execution engines* have their own limitations. Intel OAP [13] does not extensively exploit advanced code generation techniques for complex transformations and aggregations. Databricks Photon [20] and NVIDIA Spark RAPIDS [12] enable vectorized execution but lack specialized optimizations for unique feature computation patterns, such as efficient multi-window processing and custom aggregations. Similarly, Alibaba EMR [19] provides limited support for code generation in window functions, further restricting its applicability to feature extraction tasks.

**Feature Stores.** Existing online feature stores such as Feast [4] focus on rapid feature retrieval but lack mechanisms for on-the-fly, low-latency feature computation. Others, like Hopsworx [6], Tecton [27], and Feathr [5], provide limited optimization for complex feature pipelines, especially those involving multiple data streams and windowed aggregations over relative time intervals. These limitations often force users to pre-process features using separate online storage systems and to re-synchronize whenever new features are introduced, resulting in operational overhead and delayed feature availability for online ML applications.

## 3 Overview of OpenMLDB

### 3.1 System Architecture

As shown in Figure 2, OpenMLDB comprises a unified query plan generator, online real-time execution engine, offline batch execution engine, and compact time-series data and memory management.

**Unified Query Plan Generator.** To ensure consistency between offline and online feature computation, first our plan generator supports standard SQL for compatibility (e.g., MySQL queries) and SQL extensions to simplify the implementation of common features (see Section 4). Second, our plan generator uses LLVM and JIT [42, 47, 49] to transform the SQL queries into efficient machine code, where we adopt optimizations like cycle binding and compilation caching. This machine code includes direct calls to the C++ library functions shared by the offline and online execution engines.

**Online Real-Time Execution.** We developed the online engine from scratch, which is several orders of magnitude faster than

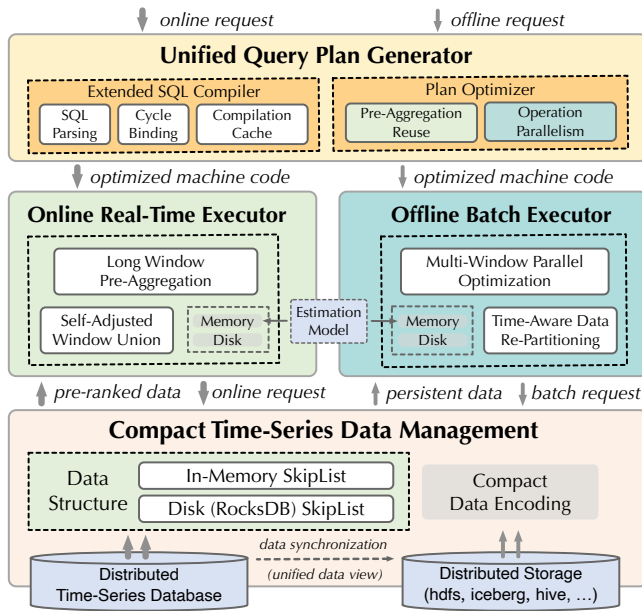


Figure 2: Architecture of OpenMLDB.

common stream processing and database systems. The underlying techniques include (1) *pre-aggregation* for enhancing functions over windows with extremely long time intervals (e.g., for years) or hotspot data and (2) *dynamic data adjusting* for efficient stream unions over multiple tables (e.g., with out-of-order stream data [65]). Note that we also adopt Zookeeper (a distributed framework [44]) to ensure high system availability and fault tolerance.

**Offline Batch Execution.** Compared to engines like Spark, the offline engine in OpenMLDB (1) improves the resource usage by parallelizing the execution of window operations over the same tables and (2) resolves the data skew problem by dynamically re-assigning window data based on the key columns and data distribution.

**Compact Time-Series Data Management.** Compared with existing systems like Spark and Redis, OpenMLDB supports a more compact data encoding format, which saves unnecessary memory consumption and enhances performance by allowing more data to be processed within limited memory space. Besides, OpenMLDB proposes an effective in-memory data structure specifically for accelerating online feature operations like retrieving one latest tuple matched with the incoming one.

Besides, to optimize memory usage, OpenMLDB offers an empirical analysis model for memory estimation, which guides memory configuration (e.g., determining the machine number and settings) and table allocation across storage engines. Besides, OpenMLDB provides a memory isolation and alerting mechanism to ensure availability when exceeding the memory limit.

### 3.2 Execution Modes

OpenMLDB supports multiple execution modes for offline and online feature computation. These modes share the same SQL syntax but impose distinct execution constraints. (1) **Offline Execution Mode:** In this mode, historical data is synchronized with the offline engine,

Table 1: Example OpenMLDB Operations (w/o in Standard SQL).

Operation	OpenMLDB
<b>Window Function</b>	<code>topN_frequency(column, top_n)</code> <code>avg_cate_where(column, condition, category)</code> <code>drawdown(column)</code> <code>ew_avg(column, smoothing factor)</code> <code>split_by_key(input_string, delimiter, kv_delimiter)</code> <code>multiclass_label(column)</code>
<b>Window Union</b>	window_clause: <code>WINDOW window_expression [ , ... ]</code> window_expression: <code>named_win AS { named_win   ( window_spec ) }</code> window_spec: <code>[ UNION ( from_item [ , ... ] ) ]</code> <code>PARTITION BY expression [ ORDER BY ... ]</code> <code>window_frame_clause [ window_attr [ , ... ] ]</code>
<b>Stream Join</b>	join_operation: <code>condition_join_operation</code> condition_join_operation: <code>from_item LAST JOIN [ ORDER BY ... ] from_item</code> join_condition: <code>join_condition</code> join_condition: <code>ON bool_expression</code>

where data scientists develop feature scripts and execute large-scale computations. Certain SQL commands (e.g., LOAD DATA) are executed asynchronously as background tasks. (2) **Online Preview Mode:** This mode enables testing newly developed feature scripts on a limited set of online data, minimizing any negative impact on online services. To ensure efficiency, it retrieves results from a data cache and constrains query complexity (e.g., limiting the number of key columns). (3) **Online Request Mode:** In this mode, each incoming request tuple is virtually inserted into the online data table. The online engine then applies the previously deployed offline script and returns a single feature result for each tuple.

## 4 Unified Query Plan Generator

### 4.1 OpenMLDB SQL

SQL is one of the most popular data processing languages. Compared to other declarative languages like Python, most relational databases like MySQL and data processing systems like Spark support SQL, making it a common choice for data processing and feature extraction in ML workflows. Additionally, the inherent constraints of SQL (e.g., execution plans without arbitrary loops) enable more targeted optimizations [31, 33].

However, standard ANSI SQL [48] has functional limitations. First, it lacks many ML feature functions (e.g., max decline percentage from a historical peak). Second, standard relational operations are inefficient in handling ML data (e.g., with ultra-high-dimension sparse features). Thus, after applying standard SQLs, users generally need to further process the output features, such as writing Python Pandas to transform into formats (e.g., LibSVM format [29]) compatible with the ML frameworks (e.g., Scikit-learn, TensorFlow). To address these problems, OpenMLDB extends standard SQL to include specialized feature computation operations<sup>3</sup>. Practical experience

<sup>3</sup>See OpenMLDB SQL syntax with over 150 built-in functions for various ML feature types at [https://openmldb.ai/docs/en/main/openmldb\\_sql/sql\\_difference.html](https://openmldb.ai/docs/en/main/openmldb_sql/sql_difference.html).

demonstrates that OpenMLDB SQL is capable of supporting feature computation in various online ML tasks [14]. Here we showcase some operations in OpenMLDB SQL (see Table 1).

**Window Function.** Compared with basic aggregations in standard SQL, OpenMLDB integrates a rich set of window functions for ML models. The following lists five typical categories:

(1) *Frequency-Based Aggregations:* The `topN_frequency` function returns the top N keys sorted by their occurrence frequency. This requires complex implementation in standard SQL, involving multiple nested queries with group and ordering operations.

(2) *Conditional Aggregations:* Similar to streaming systems that prohibit full table scans (slowing down overall efficiency), OpenMLDB extends some functions to support conditional filtering. For instance, the `avg_category_where(value, condition, category)` function computes the average of values (in string format like “shoes:20”) that meet a specified *condition*, grouped by the *category* key. Achieving similar functionality in standard SQL would require more complex constructs, involving CASE, WHERE, and ORDER BY clauses.

(3) *Time-Series Aggregations:* The `drawdown` function computes the maximum decline percentage from a historical peak to a subsequent trough, which is used in scenarios like quantitative trading to measure maximum loss. This function is not provided by standard SQL. The `ew_avg` function calculates exponentially weighted value average, which requires pre-ordered data values (see Section 7.2).

(4) *String Parsing:* The `split_by_key` function splits a string by a delimiter, treats each segment as a key-value pair, and adds each key to an output list, which standard SQL does not natively support.

(5) *Feature Signatures:* We support signature functions to mark column types by their usage in ML: (i) Label Columns are directly retained without modification. (ii) Discrete Columns are processed using a hash algorithm to generate high-dimensional feature values [60]. (iii) Continuous Columns retain their original values to produce one-dimensional dense features. In this way, we can directly output the features in LIBSVM or TFRecords format (for ML usage), avoiding to export raw table data with ultra high dimensions, such as millions of dimensions for product items.

Note these functions can be accelerated for lengthy table windows using the pre-aggregation technique (see Section 5.1).

**Window Union.** The WINDOW UNION clause aggregates data over a time window that includes data tuples from two or more tables (optimized via the self-adjusting technique in Section 5.2). This simplifies the query by avoiding complex UNION ALL operations and additional labels to distinguish which tables the tuples origin.

**Stream Join.** The LAST JOIN operation enables the matching of one most recent data tuples with the incoming data tuple (accelerated via the in-memory data structure in Section 7.2). This eliminates additional rank and filter operations in standard SQL.

With these operations, we can efficiently implement the features in Figure 1. In the SQL snippet, we first define time windows like (i) `w_union_3s` that partitions tuples by user ids (`PARTITION BY userid`), orders them by timestamp (`ORDER BY ts`) within a three-second interval before the incoming tuple (`ROWS BETWEEN 3s ...`) and (ii) long window `w_action_100d` for the interval of 100 days. Upon `w_union_3s`, we compute features like (i) the number of distinct product types (`distinct_count(action.type)`) and (ii) using

the `avg_category_where` function to conditionally compute average product prices.

```
SELECT action.*,
       distinct_count(action.type) AS product_count,
       avg_category_where(price, quantity > 1, category)
       OVER w_union_3s AS product_prices, ...
FROM actions WINDOW
w_union_3s AS (
  UNION orders PARTITION BY userid
  ORDER BY ts
  ROWS BETWEEN 3s PRECEDING AND CURRENT ROW),
w_action_100d AS ( ... ), ...;
```

## 4.2 SQL Compilation

We compile SQL into efficient C++ machine code by applying optimizations in three main aspects. (1) **Parsing Optimization:** OpenMLDB first identifies the identical column references and similar window definitions within the deployed SQL. Instead of generating separate code segments for each instance, it merges them into a unified code block. For instance, when multiple windows share the same computation template (e.g., in the form like “PARTITION BY key1 ORDER BY ts ROWS ...”), they can be merged into a single window definition, preventing redundant window construction and reducing repetitive code generation. Besides, OpenMLDB applies index optimizations to critical information (e.g., data provider) in WINDOW and LAST JOIN operations (see Section 7.2), so as to accelerate the execution process. (2) **Cyclic Binding:** When multiple aggregate functions appear in the same window, OpenMLDB applies a cyclic binding strategy. First, it compiles the minimal intermediate states required by simpler aggregates like `sum`, `min`, and `max`. These intermediate results are then reused to compute more complex aggregates such as `avg`. (3) **Compilation Cache:** OpenMLDB employs a caching mechanism that stores previously compiled code segments, allowing subsequent deployments of similar SQLs (feature scripts) to bypass the full compilation process. When a new SQL shares code patterns with an existing cached segment, OpenMLDB retrieves the relevant compiled portion and directly applies it. This reuse of compiled artifacts accelerates computation by reducing the time spent on repeated code generation and optimization.

## 5 Online Feature Computation

Beyond compiling optimization, OpenMLDB’s online execution engine offers specific optimization techniques for time-consuming operations in OpenMLDB SQL, including pre-aggregation for long window computation in Section 5.1 and self-adjusting computation for multiple window table unions in Section 5.2.

### 5.1 Long Window Pre-Aggregation

We first explain how to resolve the performance bottleneck caused by window functions involving extremely large data volumes. Such cases arise (i) when the time intervals are extremely long (e.g., spanning multiple years), which is uncommon in typical streaming scenarios, or (ii) when containing hotspot data. We refer to them as *long-window features*. On the one hand, as data volumes grow,

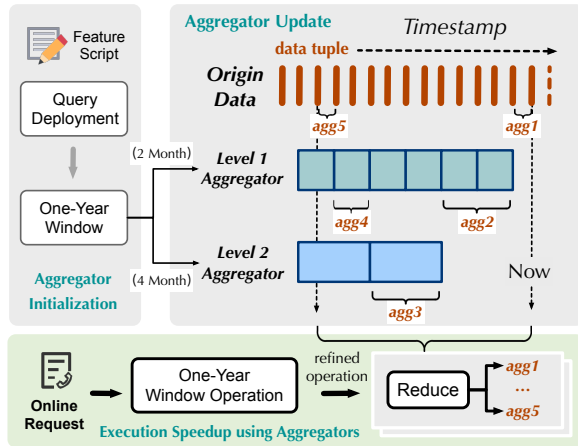


Figure 3: Long Window Pre-Aggregation.

the computation time for these long-window features increases proportionally, making it difficult to achieve real-time performance. On the other hand, processing overlapped long-window features will cause significant resource wastes.

To address these problems, we propose a long-window pre-aggregation technique. As illustrated in Figure 3, we partially compute aggregation results during data insertion. Then, for real-time requests involving *long-window features*, we simply compute by merging these pre-aggregated values, avoiding to process all raw tuples in the windows. Specifically, it includes three main steps.

**Aggregator Initialization.** We first initialize multi-level aggregators based on the long-window features in the deployed scripts. Each level corresponds to a specific time granularity and builds upon the pre-aggregated results of finer-grained levels. We determine the optimal *aggregator hierarchy* by considering factors such as query frequency, time window sizes, data distribution, and maintenance overhead. Meanwhile, we utilize data structures like segment trees [38] to efficiently manage historical computation results of these aggregators, based on which we compute statistics such as query frequencies for *aggregator hierarchy* enhancement.

**Query Refinement.** For an online request, the query optimizer identifies long-window computations and replaces expensive raw data scans with efficient lookups in the *aggregator hierarchy*. As illustrated in Figure 3, the long-window operation is transformed into the merging of (i) pre-aggregated features *agg2* and *agg4* from the daily aggregator (first level), (ii) the pre-aggregated feature *agg3* from the monthly aggregator (second level), and (iii) real-time computed features from raw data *agg1* and *agg5*. Based on the evolving query patterns, OpenMLDB can adaptively adjust the hierarchy by adding or removing aggregation levels, avoiding unnecessary computational or storage costs. For instance, to compute a feature over the past year, we can adopt aggregators at daily and monthly levels if hourly aggregators are seldom queried.

**Aggregator Update.** To maintain up-to-date aggregators without negatively impacting online data insertion, we employ an asynchronous updating strategy. First, aggregator updates are designed by assuming the “binlog\_offset” increases monotonically. To ensure this, all updates are protected within the replicator lock, which

prevents other operations, such as concurrent Put requests, from inserting conflicting updates in the middle of the sequence. By coordinating updates through the binlog, we decouple pre-aggregation computations from the main data insertion pipeline. Besides, to support failure recovery during aggregation updates with binlog, we encapsulate the aggregator update logic within update\_agg closure, ensuring atomic updates within the binlog sequence. Updates are appended to the binlog (“replicator->AppendEntry(entry, &closure)”), which not only appends the updates but also triggers asynchronous execution.

## 5.2 Self-Adjusted Window Union

Next, we explain how to optimize multi-table window unions (Window Union), which are one of the most complex operations in online feature computation. Window Union aims to match tuples from multiple stream tables over a shared time window and partitioned by common keys. For example, as shown in Figure 1, a 3-second *Window Union* ( $w_{union\_3s}$ ) returns tuples arriving within the last three seconds to the current tuple. A standard approach, used by Flink [36], employs a static key-based distribution strategy, where data tuples are routed to threads solely based on their hashed key columns. This rigid approach struggles to handle workload shifts effectively. To address this shortcoming, OpenMLDB incorporates a self-adjusting technique that dynamically refines task allocation and applies incremental computations, ensuring low response time under various workloads. It includes two main steps. **(1) On-the-Fly Load Balancing.** Our prior work [65] found that if the set of distinct keys is limited, static key-based distribution often encounters severe load imbalances. To alleviate this problem, we introduce a dynamic scheduler that periodically adjusts the mapping from keys to worker threads. By gathering runtime metrics on processing overhead, we estimate each worker thread’s load and reassign specific keys at runtime, keeping the workload even across all threads. Furthermore, multiple workers can collaborate on the same key subset if needed, improving flexibility and overall throughput. **(2) Incremental Computation.** Large windows commonly create overlapping intervals of data, making computations repetitive and resource-intensive. To avoid recalculating results from scratch, we utilize a *Subtract-and-Evict* approach [57]. As an outdated tuple leaves the window, we subtract it from the running aggregator. Then, we incorporate any newly arriving tuples incrementally.

## 6 Offline Feature Computation

### 6.1 Multi-Window Parallel Optimization

The first optimization in our offline engine focuses on *multi-window parallel optimization*, which enhances the performance of queries involving different window functions over the same table. Consider a query that involves two window functions without dependencies: (1) **Window 1** ( $w1$ ) partitions rows by “name” and orders them by “age”. (2) **Window 2** ( $w2$ ) partitions rows by “age” and orders them by “age”. In traditional data processing systems, such queries are serially computed, even when the windows have no dependency relations. This sequential processing is inefficient and presents significant optimization opportunities. OpenMLDB parallelizes these window operations by introducing two additional node types into

the execution plan: “Concat Join” and “Simple Project”. The “Concat Join” node concatenates features computed for the same data rows across different windows, marking the end of a parallel optimization segment. The “Simple Project” node represents the target columns without additional computations, marking the start of a parallel optimization segment.

Realizing parallel window computation is challenging because the order of data tuples can differ in windows using different partition keys (e.g., “name” for  $w1$  and “age” for  $w2$ ). This discrepancy makes it difficult to align and concatenate the results from parallelized windows based on their natural order. To overcome this, we introduce an *index column* to the input data, assigning a unique identifier to each data tuple. This index serves as a consistent join key across all window operations, ensuring correct alignment during concatenation regardless of the partitioning scheme.

Specifically, given the logical plan of a query, we first identify operation nodes that require parallel optimization. We perform a depth-first traversal from the root node to locate all “Concat Join” nodes inserted by the SQL parser to mark points needing parallel optimization. For each “Concat Join” node, we mark its child nodes  $\{w_i\}$  (the original window operations) to include the index column in their outputs. We also find the nearest common ancestor node of  $\{w_i\}$ , which is a “Simple Project” node. At this node, we execute a ‘Column Add’ action to insert the index column into the data stream table of these windows.

Next, with these labeled nodes, we execute the plan in a bottom-up manner. That is, each window function  $w_i$  runs independently and outputs its computed features along with the index column. After all window operations have been completed, the “Concat Join” node concatenates their results using a “Last Join” operation, with the index column as the join key (optimized via indexes). Finally, the “Concat Join” node outputs the concatenated results and removes the index column to restore the original schema.

## 6.2 Time-Aware Data Skew Resolving

The second optimization aims to resolve the data skew problem in distributed feature computation. In the distributed framework of OpenMLDB, when window computation encounters partition skew, and an imbalance is observed with excessive data in a particular partition, traditional data skew optimization methods like adding a random prefix to keys before repartitioning (known as “salting”) [30, 63] cannot be applied. Randomly prefixing keys can assign data tuples with the same group key to different partitions, leading to out-of-order processing and even inaccurate results.

To tackle this challenge, OpenMLDB introduces a data-aware parallel computation strategy that dynamically adjusts data partitions based on workload and data distribution with the following steps.

**Determine Partition Boundaries.** Given the historical data, we first evaluate the overall data to understand the distribution of the partition keys (e.g., Gender). We calculate statistical metrics to determine appropriate partition boundaries: (1) *Quantile* determines how to split the data into  $n$  equal parts. For example, setting Quantile to 4 aims to divide the data into four partitions; (2) *Percentile (PERCENTILE <sub>$i$</sub> )*: Based on the Timestamp column (e.g., used in the ORDER BY clause), we calculate boundary values that define the ranges for each partition. Data where Timestamp falls within (PERCENTILE <sub>$i$</sub> , PERCENTILE <sub>$i+1$</sub> ) belongs to the  $i$ -th partition.

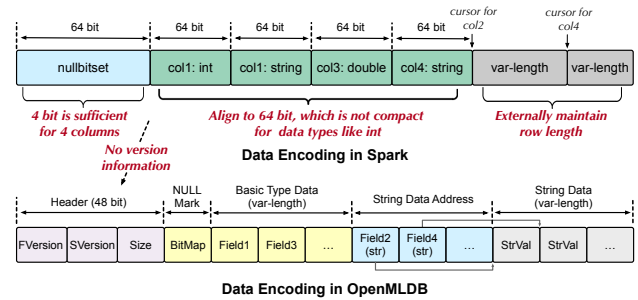


Figure 4: In-Memory Data Encoding Format.

Note we adopt HyperLogLog [40] algorithm to approximate the percentile distribution to avoid full data scan.

**Assign Repartitioning Identifiers.** Based on the distribution analysis, we tag each data tuple with: (1) *PART\_ID*: A new partition identifier that, along with the original partition key (e.g., Gender), determines the new partitions after repartitioning. (2) *EXPANDED\_ROW*: A boolean flag indicating whether the row is an original data row (false) or an expanded window row (true). In this way, we can control the repartitioning process without altering the original keys, thus preserving the continuity required for accurate window computations.

**Augment Window Data for Each Partition.** To ensure that window computations produce correct results after repartitioning (origin tuples in the same window may be of different partitions), we expand each partition’s data by including additional data tuples from preceding partitions that are necessary for computing the window functions. The process involves: (1) *Identify Required Data*: For each partition (except the first), we identify data tuples from earlier partitions that fall within the window frame needed for computations in the current partition. (2) *Mark Expanded Data*: We set the EXPANDED\_ROW flag to true for these additional data tuples to distinguish them from the original data. (3) *Union of Data*: We merge the expanded data tuples with the original partition data, ensuring each partition has all the data for accurate results.

**Redistribute Data by New Identifiers.** With the data tagged and expanded, we repartition the dataset based on both the original partition key and the new PART\_ID. This *increases the level of parallelism* by distributing the workload more evenly across multiple partitions. For example, if we split the data into four partitions, we can achieve a parallelism level of up to 8 when considering both Gender values and PART\_IDs.

**Compute on Repartitioned Data.** Finally, we perform the window computations on the repartitioned data. We compute only on data tuples where EXPANDED\_ROW is false, and the expanded data tuples (EXPANDED\_ROW is true) are utilized to provide the necessary context (e.g., identifying partition boundaries) without including them in the output.

## 7 Compact Time-Series Data Management

### 7.1 In-Memory Data Encoding

OpenMLDB uses a compact row encoding format that efficiently stores various data types while reducing memory usage. The data format combines fixed-size and variable-size data structures (see

Figure 4). Each row consists of four parts. **(1) Header (6 bytes):** The row begins with a header that encodes essential metadata, including field version, schema version, and total row size. With fewer than 64 versions, each version requires only one byte and a 32-bit value stores the row’s size. **(2) BitMap:** After the header, the BitMap indicates which columns contain NULL values. Allocated in byte units, it marks NULL fields by setting specific bits, thus avoiding the direct storage of NULL values. Its size depends on the number of columns and the presence of NULLs. **(3) Basic Type Data Fields:** Next come the basic data types (e.g., integers, floats, timestamps), stored contiguously in memory at variable lengths. For integer data, OpenMLDB uses a more compact offset calculation approach, speeding up offset determination and reducing storage costs. **(4) Variable-length Fields:** Finally, variable-length fields (e.g., strings) are stored by their offsets rather than embedding actual values. The string length is the difference between its current and previous offsets, and OpenMLDB avoids unnecessary fixed-size (32-bit) length fields. Instead, it allocates precisely as much memory as the string requires, further optimizing space utilization.

Moreover, we utilize JNI to invoke Java’s UnsafeRow data format [34], which can directly read data Rows in C, access pointers, retrieve the NullMap and each column, and obtain values through pointer offsets. In this way, we reduce the need for row encoding and decoding transformations.

**Memory Saving Example.** Consider an example row containing 20 integers, 20 floats, 20 strings (each with 1 byte), and 5 timestamps. In Spark, each row uses a NULL bit set (16 bytes), with each integer and float occupying 8 bytes, each string occupying 9 bytes (8 for data plus 1 for metadata), and each timestamp occupying 8 bytes. The total Spark row size is calculated as 556 bytes. In contrast, OpenMLDB’s row format includes a header (6 bytes), a null bit set (9 bytes), with each integer and float occupying 4 bytes, each string occupying 2 bytes (1 for data plus 1 for metadata), and each timestamp occupying 8 bytes. The total OpenMLDB row size is 255 bytes, with over 54% memory saving compared to Spark.

## 7.2 In-Memory Data Structure

For online feature computation, traditional data structures like hash tables in Redis [7] and balanced trees [32, 45] often encounter bottlenecks due to lock contention and rebalancing overhead when handling numerous simultaneous time-series insertions and complex feature computations. Redis’s hash tables also perform periodic rehashing to expand or shrink the table size, which can introduce significant latency spikes and impede real-time performance. To address these challenges, we adopt the refined skiplist data structure [50], where nodes in the first layer are sorted by key values (e.g., item ID), and each key node in the second layer points to a linked list (or a secondary skiplist) that contains all data tuples associated with that key ordered by timestamps.

**Lock-Free Concurrent Access.** We support lock-free reads and writes over the refined data skiplist (with multiple levels of linked lists that can perform insertions and deletions independently) by utilizing atomic operations for pointer updates. Specifically, when a thread inserts or deletes a node, it uses atomic compare-and-swap (CAS) instructions to update pointers, i.e., a memory location is updated only if it contains an expected value. For instance, during

insertion, the thread identifies the position where the new node should be placed and attempts to update the next pointer of the preceding node using CAS. If another thread has modified the pointer in the meantime, the CAS operation fails, and the thread retries the operation.

**Out-of-Date Data Removal.** Managing outdated data efficiently is essential for maintaining optimal performance and ensuring that OpenMLDB does not process stale information that has surpassed its time-to-live (TTL). Here, we provide an efficient outdated data removal strategy: **(1) Timestamp Ordering:** By ordering nodes based on their timestamps, the skiplist allows quick identification of outdated tuples. Since the nodes are sorted, all outdated nodes are located contiguously at the beginning or end of the list, depending on the sorting order. **(2) Batch Deletion:** To remove all data before a specific timestamp  $T$ , we traverse the skiplist from the head node and delete nodes until it reaches a node with a timestamp  $\geq T$ .

## 7.3 On-Disk Data Structure

For online data requiring persistence, we adopt RocksDB as the on-disk storage engine, used in databases like Redis [7] and TiDB [43]. First, we assign multiple indexes when creating tables for columns frequently used in operations like order-by and partition-by. Each index corresponds to a Column Family in RocksDB. Different column families have separate Sorted String Table (SST) files and individual data eviction policies, but they share a single Memory Table (memtable). This means that while each Column Family manages its on-disk storage and data eviction independently, they all write to the same in-memory data structure before flushing data to disk. In our implementation, we utilize the refined data skiplist as the Memory Table, where a key and a timestamp (ts) are concatenated to form a composite key. The RocksDB keys are pre-sorted, so data with the same key is grouped together, facilitating data queries over a specified time range. We support data eviction in RocksDB by parsing the keys and timestamps in the skiplist and removing data whose timestamp is out-of-date.

## 8 Memory Management Mechanisms

### 8.1 Table-Level Memory Configuration

OpenMLDB workloads often require substantial memory resources, particularly when dealing with large-scale feature computations. To guide users in choosing appropriately sized memory configurations, we provide an empirical memory estimation model.

**Memory Usage Estimation.** Inspired by established practices in feature computation [15], we use the following model to estimate total memory usage:

$$\text{mem\_total} = \sum_{i=1}^{n_{\text{table}}} n_{\text{replica}_i} \left[ \left( \sum_{j=1}^{n_{\text{index}}} n_{\text{pk}_{ij}} \times (|pk_{ij}| + 156) \right) + n_{\text{index}_i} \times n_{\text{row}_i} \times C + K \times n_{\text{row}_i} \times |\text{row}_i| \right]$$

where  $n_{\text{table}}$  is the total number of tables, and  $n_{\text{index}}$  is the number of indexes per table. For each table  $i$ ,  $n_{\text{replica}_i}$  is the number of replicas (full table data copies),  $n_{\text{pk}_{ij}}$  is the number of unique primary keys on the  $j$ th index column, and  $|pk_{ij}|$  is the average key column length. The parameters  $n_{\text{index}_i}$  and  $n_{\text{row}_i}$  denote the number of

indexes and tuples in table  $i$ , respectively, while  $|row_i|$  is the average tuple length. The factor  $C$  depends on the table type: set  $C = 70$  for “latest” (storing recent data for a given key) or “absorlat” (with additional logic like combining recent entries) tables, and  $C = 74$  for “absolute” (data keyed by a timestamp or an absolute reference) or “absandlat” (e.g., data accessible by absolute timestamps) tables. The parameter  $K$ , representing the number of data copies stored, varies between 1 and  $n_{index}$ . For instance, suppose we have a “latest” table with 1 million rows, average row length 300 bytes, two indexes, and replicas equal 2. Each primary key is 16 bytes. Using  $C = 70$  and  $K = 1$ , the memory usage is about 1.568 GB. Note memory consumption by pre-aggregation are typically negligible compared to the main memory consumption and so is not reflected here.

With the estimated memory usage, users can flexibly assign storage engines to tables in OpenMLDB. For instance, if the estimated memory usage for a table fits within available memory resources and ultra-low latency (around 10 milliseconds) is required, the in-memory storage engine is recommended. If memory resources are limited or the estimated memory usage exceeds available memory, and a latency of 20-30 milliseconds is acceptable, the disk-based storage engine may be considered, leading to approximately 80% savings in hardware costs (see Section 9.2).

## 8.2 Runtime Memory Management

Moreover, to ensure high stability and prevent memory exhaustion, which can lead to tablets being terminated by the operating system and result in service downtime, OpenMLDB introduces two key features: (1) **Memory Resource Isolation** provides a tablet-level configuration parameter (`max_memory_mb`) that limits the maximum memory usage of a tablet. When memory usage exceeds this limit, write operations fail but read operations continue, keeping the service online. Users can address memory shortages through scaling, shard migration, or other methods with minimal impact on production systems. (2) **Memory Alerting Mechanism**: An alerting mechanism notifies relevant personnel or systems when memory usage surpasses a preset threshold.

## 9 Evaluation

### 9.1 Experimental Setup

**Workloads.** We evaluate OpenMLDB using a micro-benchmark (MicroBench) and various real-world production workloads.

• **MicroBench.** We implemented a Java-based testing tool employing the OpenMLDB Java SDK and Jedis to compare performance and memory consumption. Due to differences in supported data types and storage layouts, the data insertion procedure for each system was slightly adjusted. To reflect realistic usage, we designed time-series-based tests (with three stream tables) with multiple adjustable parameters (e.g., number of windows, number of LAST JOIN operations).

• **TalkingData.** The open TalkingData dataset (covering around 200 million clicks over four days [22]) contains common column types (e.g., strings, number values, timestamps). Since the dataset is time-series-based and many data tuples share the same ip key, we created a table indexed by the ip column, allowing us to calculate the memory usage savings by OpenMLDB and the baselines.




Figure 5: Online MicroBench Performance Comparison.

- **RTP.** The RTP service at Akulaku is an item-ranking framework with a data warehouse exceeding 1,000 TB. It supports various online machine learning models, such as Logistic Regression (LR) and Gradient Boosted Decision Trees (GBDT), as well as diverse feature formats. This framework powers multiple online ML tasks, including user risk ranking, device fingerprinting, and facial recognition.
- **GLQ.** The GLQ workload consists of 10 billion geospatial tuples linked to GPS coordinates. It involves complex, full-scale geospatial queries (e.g., proximity searches, route optimization, clustering), which are extremely resource-intensive and can cause out-of-memory (OOM) errors in traditional systems.

**Test Settings.** For MicroBench, we use four servers of identical configurations: each with 512GB of memory, dual Intel Xeon E5-2630 v4 CPUs (2.20GHz), a 1Gbps network card, and running CentOS-7. Three servers run the OpenMLDB service, and the fourth hosts the client application using the OpenMLDB Java SDK for testing. For other workloads, experiments are conducted on 16 servers of similar configurations.

## 9.2 Overall Performance

### 9.2.1 Online Feature Computation

**MicroBench Performance.** We first compare OpenMLDB with three commonly used baselines for in-memory online analytics: (1) Trino+Redis: We pair Redis (a popular in-memory data store [7]) with Trino (an ANSI SQL-compliant engine [37]) to handle SQL queries over Redis data. (2) MySQL (in-mem): We test MySQL configured with the MEMORY storage engine to compute online features over the stored relational data [24]. (3) DuckDB: We also compare DuckDB, an embedded analytical database [51] designed for online analytical processing. As shown in Figure 5, OpenMLDB achieves significant performance improvement in both latency and throughput compared to all the baselines on MicroBench. In terms of latency, OpenMLDB outperforms MySQL (in-mem) by over 68.4%, DuckDB by 87.7%, and Trino+Redis by over 96%, respectively. OpenMLDB also achieves throughput gains by over 17 times higher than the baselines. The reasons are three-fold.

• OpenMLDB employs a C++-based compilation framework that specializes window operations across multiple granularities (e.g., daily, weekly, and monthly) in a single integrated pass. OpenMLDB also uses LLVM-based Just-In-Time (JIT) compilation to fuse and streamline aggregate functions (e.g., sum and avg). Instead, MySQL (in-mem) relies heavily on interpreted SQL execution, and DuckDB does not incorporate aggressive cross-operation optimizations. Moreover, compared to Trino+Redis’s Java framework with frequent RPC calls, OpenMLDB benefits from more direct control over low-level




Figure 6: RTP Performance Comparison.

operation optimizations and can apply advanced caching strategies that further reduce overheads for reusable feature scripts.

- OpenMLDB leverages the in-memory data structure to pre-rank tuples by timestamps. As a result, data retrieval and scanning occur in a nearly linear time complexity with respect to input sizes. In contrast, MySQL (in-mem) often incurs overhead from disk-based indexing, while DuckDB, though in-memory and columnar, may still require additional passes for complex temporal queries. Trino+Redis, on the other hand, spreads window-state management over multiple operators, introducing extra latency for state lookups and aggregation steps.

- OpenMLDB dynamically redistributes data by keys and employs incremental computation strategies for windowed operations. Instead of reprocessing entire datasets for each new computation like MySQL (in-mem) and DuckDB, OpenMLDB updates only the necessary segments of the data stream. Likewise, compared to Trino+Redis’s distributed streaming model, which often involves shuffling and checkpointing overhead, OpenMLDB’s key-based redistribution and incremental approaches minimize redundant computations, thereby significantly reducing overall processing time and improving real-time responsiveness.

**RTP Performance.** The Akulaku company has integrated OpenMLDB into its ML pipeline to streamline real-time feature generation and ranking [14]. Figure 6 shows the real performance tested in Akulaku’s RTP business, which compares OpenMLDB with GreenPlum (a typical MPP database [8]) and the streaming system Flink [36]. We can see OpenMLDB significantly outperforms the baselines for real-time TopN queries. GreenPlum incurs prohibitive recomputations for new data tuples, causing substantial delays. Flink, although good at processing basic time windows, is also not well optimized for TopN ranking. In contrast, OpenMLDB *scales nearly linearly*, achieving Top1 queries in about 0.98 ms and Top8 queries around 5 ms, significantly better than Flink’s sub-100 ms range. The improvement is due to (1) OpenMLDB’s runtime dynamic compilation and LLVM IR code generation, which embed query parameters directly into machine code to improve execution efficiency without sacrificing generality, and (2) OpenMLDB’s in-memory data structure pre-ranks stream data by keys and timestamps, thereby minimizing runtime sorting overhead.

**Memory Consumption.** We compare memory usage between OpenMLDB and Trino+Redis. As shown in Table 2, we can see OpenMLDB achieves over 50% memory saving than Trino+Redis, which can be extremely vital for practical usage and saving machine costs. Specifically, for smaller datasets (10K tuples), OpenMLDB uses 74.77% less memory than Trino+Redis. As dataset size grows to

Table 2: Memory resource saved by OpenMLDB (bytes).

#-Tuples	RedisMem	OpenMLDB Mem	Reduction
10,000	9,272,328	2,339,699	↓ 74.77%
100,000	48,501,288	15,624,290	↓ 67.79%
1,000,000	215,323,024	105,722,441	↓ 50.90%
10,000,000	1,897,343,984	1,008,276,458	↓ 46.86%
184,903,890	34,071,049,864	18,513,271,540	↓ 45.66%




Figure 7: Offline MicroBench Performance Comparison.

hundreds of millions of tuples, the memory reduction stabilizes around 45.66%. This memory efficiency stems from OpenMLDB’s compact time-series encoding and specialized indexing, reducing overhead from repeated keys and non-compact data layouts typical in Trino+Redis or other general-purpose stores. By minimizing memory usage, OpenMLDB not only lowers hardware costs but also processes more data in-memory, enabling more scalable and lower-latency feature computations than baseline systems that must either add more hardware or spill to disk.

## 9.2.2 Offline Feature Computation

**MicroBench Performance.** As shown in Figure 7, on the offline MicroBench, OpenMLDB outperforms Spark significantly, delivering a 2.6x speedup on single-window queries, a 6.3x speedup on multiple-window workloads. Against skewed data (imbalances across the distributed data partitions), OpenMLDB reduces latency by 7.2x with skew optimization (180s vs. 1302s). The reasons are three-fold. First, OpenMLDB employs C++ execution libraries and leverages just-in-time (JIT) compilation to generate efficient native machine code. In contrast, Spark’s query execution relies heavily on the JVM runtime, resulting in additional overhead from garbage collection and interpretation. Second, OpenMLDB prioritizes direct in-memory window operations to reduce CPU cache misses. It efficiently parallelizes time-window computations and streamlines memory usage. While Spark must orchestrate multiple stages and shuffles, incurring expensive serialization, deserialization, and data movement. Third, OpenMLDB removes unnecessary initialization routines and repetitive object constructions. By trimming down code paths and utilizing a more straightforward execution model, OpenMLDB enhances both performance and maintainability. In contrast, Spark’s more complex execution pipelines and object-heavy data structures introduce additional execution overhead.

**GLQ Performance.** Moreover, we compare the performance of OpenMLDB and Spark in handling real geographic location scenarios involving complex queries. These queries necessitate evaluating the relative relationships among all GPS coordinates, thus requiring computation over the entire dataset for each query. This





Figure 8: Offline GLQ Performance Comparison.



(a) Latency (b) Throughput

Figure 9: Pre-Aggregation Performance.


makes traditional SQL-based expressions inadequate and demands OpenMLDB SQL. As shown in Figure 8, OpenMLDB achieves near 30-millisecond response time, with execution speedups ranging from approximately 5x to over 22x compared to Spark as the hyper-parameter  $N$  increases from 7 to 10. By leveraging LLVM-based compilation optimizations, OpenMLDB generates more efficient execution code than Spark. Additionally, Spark’s memory structure exhibits inefficiencies in handling homogeneous, single-purpose tasks and encounters out-of-memory (OOM) errors for full-table queries. These observations underscore OpenMLDB’s superiority in optimizing data formats and memory management.

### 9.3 Ablation Study

#### 9.3.1 Long Window Pre-Aggregation

To evaluate the impact of pre-aggregation on long windows, we configure a series of SQL queries with varying window sizes, ranging from 100K to 5000K tuples. As shown in Figure 9, the performance improvement from enabling pre-aggregation is substantial. Without pre-aggregation, as the window size grows, OpenMLDB must repeatedly scan and aggregate large volumes of raw data, leading to exponential growth in latency and a precipitous drop in throughput. In contrast, with pre-aggregation enabled, intermediate partial aggregates are precomputed and cached, drastically reducing the overhead for subsequent queries over the same window.

Furthermore, we evaluate the effectiveness of long window pre-aggregation over a larger dataset of 860,000 tuples. As shown in Figure 10, we observe that the long window pre-aggregation in OpenMLDB significantly improves the efficiency of a request with little SQL modification and a slightly higher data loading overhead. Specifically, when specifying the usage of the long window pre-aggregation, i.e., adding `deploy test1 OPTIONS(long_windows="w1:1d")` in the SQL, the response time for a request obtains a 45x reduction (i.e., from 300ms to 6ms) with this optimization technique. Therefore, while pre-aggregation introduces a small overhead during data loading, OpenMLDB’s online computation efficiency can be greatly improved, especially for queries with repeated window aggregations.



(a) Long Window SQL (b) Latency Comparison

Figure 10: Performance of Long-Window Optimization: (a) Testing Feature Script and (b) Latency Comparison.




Figure 11: Multi-Window Parallel Optimization Performance.

#### 9.3.2 Multi-Table Window Union Optimization

Next, we evaluate the performance of multi-table window union operations. Without the self-adjusting techniques introduced in Section 5.2, throughput severely degrades as window sizes increase. For example, the static execution approach (e.g., the one in Apache Flink [36]) has a throughput of around 1000 tuples per second when the window size is 10,000 due to the need to reprocess extensive overlapping data intervals and handle skewed key distributions inefficiently. In contrast, OpenMLDB maintains stable, high throughput (approximately 1,000,000 tuples/sec) across varying window sizes by leveraging its self-adjusting load balancing and incremental computation strategies. The dynamic key-to-thread mapping continuously redistributes workload to address shifting access patterns, mitigating hot spots that the static method struggles with.

#### 9.3.3 Multi-Window Parallel Optimization

As shown in Figure 11, we observe that OpenMLDB with parallel optimization yields significant performance improvements. With small windows containing 1,000 rows, the performance of the smaller window becomes completely overshadowed in OpenMLDB with parallel optimization. The overall task time perceived by the user is primarily the duration of the window with the longest execution time, resulting in a 4.8 times performance increase. In medium-sized windows, performance improves by 5.3 times, and parallel optimization on large windows yields a 4.6 times improvement compared to Spark. By concurrently processing multiple windows, OpenMLDB spreads workloads across CPU cores more evenly and avoids serialization bottlenecks. Note while this parallel optimization in OpenMLDB introduces overhead of some additional merging operations, we have minimized it through optimizations like the Last Join syntax for one-to-one correspondence with higher performance than traditional left joins (see Section 4.1).




Figure 12: Data Skew Optimization Performance

### 9.3.4 Data Skew Optimization

Figure 12 shows that, after applying data skew optimization, the performance has significantly improved compared to Spark and OpenMLDB without skew optimization. Although OpenMLDB without skew optimization already achieves 4 times performance improvement over Spark, the improvement becomes even greater with skew optimization, such as skew 2 (the default value indicates the doubled data partition number). At skew 4, performance is 10.1 times higher over Spark and more than 2 times over without optimization. This improvement stems from the data-aware strategy that adjusts distributed partitions based on workload and data distributions.

## 9.4 Evaluation on Hyper-Parameters

Furthermore, we investigate how various parameters affect OpenMLDB’s performance, including (1) *System parameters*: number of threads, whether to enable pre-aggregation optimization, (2) *Query parameters*: number of windows, number of LAST JOINS, (3) *Data parameters*: number of data tuples within a window, cardinality of the indexed column (i.e., the number of unique values after deduplication). Pre-aggregation optimization is usually only meaningful when the data volume within the window is huge (e.g., millions of data tuples), so it is turned off by default.

Table 3: Performance for Different Feature Numbers.

#-Column	#-Feature	TP50	TP90	TP95	TP99	TP999
10	20	0.6	0.8	0.8	1.0	1.9
100	210	2.0	2.8	2.5	4.4	6.6
1000	2100	11.7	14.7	15.9	19.8	44.8

**Number of Features.** To evaluate the scalability with respect to feature dimensions, we consider scenarios with varying numbers of columns and derived features. For instance, increasing from 10 columns (20 features) to 1,000 columns (2,100 features) leads to higher latency. However, the latency remains within tens of milliseconds, which is still acceptable for most online analytical workloads. As shown in Table 3, even at high feature counts, the observed latency is modest, confirming OpenMLDB’s ability to handle complex feature pipelines without significant performance degradation.

**Number of Threads.** Results show increasing threads improves throughput but can slightly increase latency. Beyond 20 threads, latency grows but remains within single-digit milliseconds. This trade-off allows users to tune resource usage for their latency vs. throughput requirements. Such scaling behavior is more predictable and manageable than in Spark or MySQL, where concurrency often introduces non-linear slowdowns or GC overhead.




Figure 13: Performance under Different Window Numbers.




Figure 14: Performance under Different Join Numbers.

**Number of Windows and Data Volume.** Figure 13 demonstrates that, as the number of windows or data size per window grows, the latency of OpenMLDB increases modestly but generally stays under 10 ms. While throughput decreases, OpenMLDB still maintains practical performance for complex analytics.

**Number of Joins.** Figure 14 shows that feature scripts with more LAST JOIN operations slightly increase latency (still under 5ms) and reduce throughput marginally, remaining above 6k QPS. This is critical for feature pipelines that require joining multiple reference tables. Other methods like MySQL or DuckDB would see far larger latency spikes due to repeated index lookups and lack of incremental computation (details in Section 9.2).

## 10 Conclusion

In this paper, we presented OpenMLDB, an industrial-grade feature computation system. OpenMLDB unifies online and offline feature computation through a unified query plan generator and applies compilation-level optimizations like cycle binding and code caching. For online feature computation, OpenMLDB uses pre-aggregation and dynamic data adjustments to efficiently handle long windows and multi-table unions. For offline tasks, OpenMLDB enables multi-window parallelism and time-aware data repartitioning to scale historical feature computation. We also designed compact in-memory formats and stream-based structures to speed up time-series data access. Extensive testing and real-world deployment show that OpenMLDB achieves faster execution, better resource efficiency, and stable performance across different workloads.

## Acknowledgments

This paper was supported by National Key R&D Program of China (2023YFB4503600), NSF of China (62232009, 62102215), Shenzhen Project (CJGJZD20230724093403007), Zhongguancun Lab, and Beijing National Research Center for Information Science and Technology (BNRist). Bingsheng He is the corresponding author.

## References

- [1] [n.d.]. <https://cassandra.apache.org>. Last accessed on 2024-12.
- [2] [n.d.]. <https://doris.apache.org>. Last accessed on 2024-12.
- [3] [n.d.]. <https://druid.apache.org>. Last accessed on 2024-12.
- [4] [n.d.]. <https://github.com/feast-dev/feast>. Last accessed on 2024-11.
- [5] [n.d.]. <https://github.com/feathr-ai/feathr>. Last accessed on 2024-11.
- [6] [n.d.]. <https://github.com/logicalclocks/hopsworks>. Last accessed on 2024-11.
- [7] [n.d.]. <https://github.com/redis/redis>. Last accessed on 2024-11.
- [8] [n.d.]. <https://greenplum.org>. Last accessed on 2024-12.
- [9] [n.d.]. <https://hbase.apache.org>. Last accessed on 2024-12.
- [10] [n.d.]. <https://mariadb.org>. Last accessed on 2024-12.
- [11] [n.d.]. <https://medium.com/engineering-varo/feature-store-challenges-and-considerations-d1d59c070634>. Last accessed on 2024-12.
- [12] [n.d.]. <https://nvidia.github.io/spark-rapids>. Last accessed on 2024-12.
- [13] [n.d.]. <https://oap-project.github.io/latest>. Last accessed on 2024-11.
- [14] [n.d.]. <https://openmldb.ai/en/akulaku-real-time-feature-extraction-for-ai-powered-risk-control>. Last accessed on 2024-12.
- [15] [n.d.]. <https://openmldb.ai/en/blog/>. Last accessed on 2024-12.
- [16] [n.d.]. <https://openmldb.ai/en/the-commercial-bank-driving-the-ai-based-anti-fraud-with-efficient-feature-extraction/>. Last accessed on 2024-12.
- [17] [n.d.]. <https://opentsdb.net>. Last accessed on 2024-12.
- [18] [n.d.]. <https://tdengine.com>. Last accessed on 2024-12.
- [19] [n.d.]. <https://www.alibabacloud.com/product/emapreduce>. Last accessed on 2024-12.
- [20] [n.d.]. <https://www.databricks.com/product/photon>. Last accessed on 2024-11.
- [21] [n.d.]. <https://www.influxdata.com>. Last accessed on 2024-12.
- [22] [n.d.]. <https://www.kaggle.com/competitions/talkingdata-adtracking-fraud-detection>. Last accessed on 2024-12.
- [23] [n.d.]. <https://www.mongodb.com>. Last accessed on 2024-12.
- [24] [n.d.]. <https://www.mysql.com>. Last accessed on 2024-12.
- [25] [n.d.]. <https://www.postgresql.org>. Last accessed on 2024-12.
- [26] [n.d.]. <https://www.starrocks.io>. Last accessed on 2024-12.
- [27] [n.d.]. <https://www.tecton.ai>. Last accessed on 2024-12.
- [28] 2018. Forecast: The Business Value of Artificial Intelligence. In *Gartner*. <https://www.gartner.com/en/documents/3868267>
- [29] Abdiansah Abdiansah and Retantyo Wardoyo. 2015. Time complexity analysis of support vector machines (SVM) in LibSVM. *Int. J. Comput. Appl* 128, 3 (2015), 28–34.
- [30] Adrian-Ioan Argesanu and Gheorghe-Daniel Andreescu. 2021. Accelerating Data Ingress for Range-Scan Optimized HBase Instances. In *15th IEEE International Symposium on Applied Computational Intelligence and Informatics, SACI 2021, Timisoara, Romania, May 19-21, 2021*. IEEE, 341–344. <https://doi.org/10.1109/SACI1354.2021.9465633>
- [31] Michael Armbrust, Reynold S Xin, Cheng Lian, Yin Huai, Davies Liu, Joseph K Bradley, Xiangrui Meng, Tomer Kaftan, Michael J Franklin, Ali Ghodsi, et al. 2015. Spark sql: Relational data processing in spark. In *Proceedings of the 2015 ACM SIGMOD international conference on management of data*. 1383–1394.
- [32] Tiemo Bang, Ismail Oukid, Norman May, Ilia Petrov, and Carsten Binnig. 2020. Robust Performance of Main Memory Data Structures by Configuration. In *Proceedings of the 2020 International Conference on Management of Data, SIGMOD Conference 2020, online conference [Portland, OR, USA], June 14-19, 2020*. ACM, 1651–1666. <https://doi.org/10.1145/3318464.3389725>
- [33] Edmon Begoli, Jesús Camacho-Rodríguez, Julian Hyde, Michael J Mior, and Daniel Lemire. 2018. Apache calcite: A foundational framework for optimized query processing over heterogeneous data sources. In *Proceedings of the 2018 International Conference on Management of Data*. 221–230.
- [34] Ionut Boicu. 2019. *Adaptive on-the-fly compressed execution in spark*. Ph.D. Dissertation. Master's Thesis. Vrije Universiteit Amsterdam.
- [35] Jie Cai, Jiawei Luo, Shulin Wang, and Sheng Yang. 2018. Feature selection in machine learning: A new perspective. *Neurocomputing* 300 (2018), 70–79. <https://doi.org/10.1016/j.neucom.2017.11.077>
- [36] Paris Carbone, Asterios Katsifodimos, Stephan Ewen, Volker Markl, Seif Haridi, and Kostas Tzoumas. 2015. Apache Flink™: Stream and Batch Processing in a Single Engine. *IEEE Data Eng. Bull.* 38, 4 (2015), 28–38. <http://sites.computer.org/debull/A15dec/p28.pdf>
- [37] Cristian Cardas, José F. Aldana-Martín, Antonio M. Burgueño-Romero, Antonio J. Nebro, Jose Manuel Mateos, and Juan Jesús Sánchez Sánchez. 2023. On the performance of SQL scalable systems on Kubernetes: a comparative study. *Clust. Comput.* 26, 3 (2023), 1935–1947. <https://doi.org/10.1007/S10586-022-03718-9>
- [38] Mark de Berg, Marc van Kreveld, Mark Overmars, Otfried Cheong Schwarzkopf, Mark de Berg, Marc van Kreveld, Mark Overmars, and Otfried Cheong Schwarzkopf. 2000. More geometric data structures: Windowing. *Computational Geometry: algorithms and applications* (2000), 211–233.
- [39] Jeffrey Dean and Sanjay Ghemawat. 2008. MapReduce: simplified data processing on large clusters. *Commun. ACM* 51, 1 (2008), 107–113.
- [40] Philippe Flajolet, Éric Fusy, Olivier Gandouet, and Frédéric Meunier. 2007. Hyperloglog: the analysis of a near-optimal cardinality estimation algorithm. *Discrete mathematics & theoretical computer science* Proceedings (2007).
- [41] Gartner. [n.d.]. The Logical Feature Store: Data Management for Machine Learning. Last accessed on 2024-02.
- [42] Michal Gregor and Juraj Spalek. 2017. Using LLVM-based JIT Compilation in Genetic Programming. *CoRR* abs/1701.05730 (2017). [arXiv:1701.05730](http://arxiv.org/abs/1701.05730) <http://arxiv.org/abs/1701.05730>
- [43] Dongxu Huang, Qi Liu, Qiu Cui, Zhuhe Fang, Xiaoyu Ma, Fei Xu, Li Shen, Liu Tang, Yuxing Zhou, Menglong Xue, et al. 2020. TiDB: a Raft-based HTAP database. *Proceedings of the VLDB Endowment* 13, 12 (2020), 3072–3084.
- [44] Patrick Hunt, Mahadev Konar, Flavio Paiva Junqueira, and Benjamin C. Reed. 2010. ZooKeeper: Wait-free Coordination for Internet-scale Systems. In *Proceedings of the 2010 USENIX Annual Technical Conference, USENIX ATC 2010, Boston, MA, USA, June 23-25, 2010*. USENIX Association. <https://www.usenix.org/conference/usenix-atc-10/zookeeper-wait-free-coordination-internet-scale-systems>
- [45] Hongjun Lu, Yuet Yeung Ng, and Zengping Tian. 2000. T-Tree or B-Tree: Main Memory Database Index Structure Revisited. In *Australasian Database Conference, ADC 2000, Canberra, Australia, January 31 - February 3, 2000*. IEEE Computer Society, 65–73. <https://doi.org/10.1109/ADC.2000.819815>
- [46] Terry J. Lyons and Andrew D. McLeod. 2022. Signature Methods in Machine Learning. *CoRR* abs/2206.14674 (2022). <https://doi.org/10.48550/ARXIV.2206.14674>
- [47] Miao Ma, Zhengyi Yang, Kongzhang Hao, Liuyi Chen, Chunling Wang, and Yi Jin. 2023. An Empirical Analysis of Just-in-Time Compilation in Modern Databases. In *Databases Theory and Applications - 34th Australasian Database Conference, ADC 2023, Melbourne, VIC, Australia, November 1-3, 2023, Proceedings (Lecture Notes in Computer Science)*, Zhifeng Bao, Renata Borovica-Gajic, Ruihong Qiu, Farhana Murtaza Choudhury, and Zhengyi Yang (Eds.), Vol. 14386. Springer, 227–240. [https://doi.org/10.1007/978-3-031-47843-7\\_16](https://doi.org/10.1007/978-3-031-47843-7_16)
- [48] Jim Melton. 1996. SQL Language Summary. *ACM Comput. Surv.* 28, 1 (1996), 141–143. <https://doi.org/10.1145/234313.234374>
- [49] Vaivasswatha Nagaraj, Jacob Johannsen, Anton Trunov, George Pirlea, Amrit Kumar, and Ilya Sergey. 2020. Compiling a Higher-Order Smart Contract Language to LLVM. *CoRR* abs/2008.05555 (2020). [arXiv:2008.05555](https://arxiv.org/abs/2008.05555) <https://arxiv.org/abs/2008.05555>
- [50] William W. Pugh. 1990. Skip Lists: A Probabilistic Alternative to Balanced Trees. *Commun. ACM* 33, 6 (1990), 668–676. <https://doi.org/10.1145/78973.78977>
- [51] Mark Raasveldt and Hannes Mühleisen. 2019. Duckdb: an embeddable analytical database. In *Proceedings of the 2019 International Conference on Management of Data*. 1981–1984.
- [52] Salman Salloum, Ruslan Dautov, Xiaojun Chen, Patrick Xiaogang Peng, and Joshua Zhexue Huang. 2016. Big data analytics on Apache Spark. *International Journal of Data Science and Analytics* 1 (2016), 145–164.
- [53] Robert Schulze, Tom Schreiber, Ilya Yatsishin, Ryadh Dahimene, and Alexey Milovidov. 2024. ClickHouse-Lightning Fast Analytics for Everyone. *Proceedings of the VLDB Endowment* 17, 12 (2024), 3731–3744.
- [54] Swaminathan Sivasubramanian. 2012. Amazon dynamoDB: a seamlessly scalable non-relational database service. In *Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data*. 729–730.
- [55] Ji Sun and Guoliang Li. 2019. An End-to-End Learning-based Cost Estimator. *Proc. VLDB Endow.* 13, 3 (2019), 307–319. <https://doi.org/10.14778/3368289.3368296>
- [56] Rebecca Taft, Irfan Sharif, Andrei Matei, Nathan VanBenschoten, Jordan Lewis, Tobias Grieger, Kai Niemi, Andy Woods, Anne Birzin, Raphael Poss, et al. 2020. Cockroachdb: The resilient geo-distributed sql database. In *Proceedings of the 2020 ACM SIGMOD international conference on management of data*. 1493–1509.
- [57] Kanat Tangwongsan, Martin Hirzel, and Scott Schneider. 2017. Low-Latency Sliding-Window Aggregation in Worst-Case Constant Time (*DEBS '17*). Association for Computing Machinery, New York, NY, USA, 66–77. <https://doi.org/10.1145/3093742.3093925>
- [58] Ashish Thusoo, Joydeep Sen Sarma, Namit Jain, Zheng Shao, Prasad Chakka, Suresh Anthony, Hao Liu, Pete Wyckoff, and Raghotham Murthy. 2009. Hive: a warehousing solution over a map-reduce framework. *Proceedings of the VLDB Endowment* 2, 2 (2009), 1626–1629.
- [59] Jiayi Wang, Chengliang Chai, Jiabin Liu, and Guoliang Li. 2021. FACE: A Normalizing Flow based Cardinality Estimator. *VLDB* 15, 1 (2021), 72–84. <http://www.vldb.org/pvldb/vol15/p72-li.pdf>
- [60] Kilian Weinberger, Anirban Dasgupta, John Langford, Alex Smola, and Josh Attenberg. 2009. Feature hashing for large scale multitask learning. In *Proceedings of the 26th annual international conference on machine learning*. 1113–1120.
- [61] Zhenkun Yang, Chuanhui Yang, Fusheng Han, Mingqiang Zhuang, Bing Yang, Zhifeng Yang, Xiaojun Cheng, Yuzhong Zhao, Wenhui Shi, Huafeng Xi, et al. 2022. OceanBase: a 707 million tpmC distributed relational database system. *Proceedings of the VLDB Endowment* 15, 12 (2022), 3385–3397.
- [62] Matei Zaharia, Reynold S Xin, Patrick Wendell, Tathagata Das, Michael Armbrust, Ankur Dave, Xiangrui Meng, Josh Rosen, Shivaram Venkataraman, Michael J Franklin, et al. 2016. Apache spark: a unified engine for big data processing. *Commun. ACM* 59, 11 (2016), 56–65.
- [63] Eftim Zdravevski, Petre Lameski, and Andrea Kulakov. 2016. Row Key Designs of NoSQL Database Tables and Their Impact on Write Performance. In *24th*

- Euromicro International Conference on Parallel, Distributed, and Network-Based Processing, PDP 2016, Heraklion, Crete, Greece, February 17-19, 2016*. IEEE Computer Society, 10–17. <https://doi.org/10.1109/PDP.2016.84>
- [64] Hao Zhang, Xianzhi Zeng, Shuhao Zhang, Xinyi Liu, Mian Lu, and Zhao Zheng. 2023. Scalable Online Interval Join on Modern Multicore Processors in OpenMLDB. In *39th IEEE International Conference on Data Engineering, ICDE 2023, Anaheim, CA, USA, April 3-7, 2023*. IEEE, 3031–3042. <https://doi.org/10.1109/ICDE55515.2023.00232>
- [65] Hao Zhang, Xianzhi Zeng, Shuhao Zhang, Xinyi Liu, Mian Lu, and Zhao Zheng. 2023. Scalable Online Interval Join on Modern Multicore Processors in OpenMLDB. In *2023 IEEE 39th International Conference on Data Engineering (ICDE)*. 3031–3042. <https://doi.org/10.1109/ICDE55515.2023.00232>
- [66] Xuanhe Zhou, Cheng Chen, Kunyi Li, Bingsheng He, Mian Lu, Qiaosheng Liu, Wei Huang, Guoliang Li, Zhao Zheng, and Yuqiang Chen. 2023. FEBench: A Benchmark for Real-Time Relational Data Feature Extraction. *Proc. VLDB Endow* 16, 12 (2023), 3597–3609. <https://doi.org/10.14778/3611540.3611550>