


# Relation-Based In-Database Stream Processing

**Christian Winter**, Thomas Neumann, Alfons Kemper

Technical University of Munich


CDMS @ VLDB 2023




*Uhrenturm der TUM*

→ Stream data analytics is growing increasingly complex


- Stream data analytics is growing increasingly complex
- Streams are enriched and analyzed with historic data
  - Business reporting
  - Log monitoring
  - Logistics



- Stream data analytics is growing increasingly complex
- Streams are enriched and analyzed with historic data
  - Business reporting
  - Log monitoring
  - Logistics
- ⚠ Stream processing engines inefficient for historic data




- Stream data analytics is growing increasingly complex
- Streams are enriched and analyzed with historic data
  - Business reporting
  - Log monitoring
  - Logistics
- ⚠ Stream processing engines inefficient for historic data
- ⚠ Database systems not built for ephemeral data




→ **Complex stream analytics in a single system is hard**


- Historic data in stream processing engines
  - ⚠ Read-only
  - ⚠ Optimized for filter-and-aggregate queries



- Historic data in stream processing engines
  - ⚠ Read-only
  - ⚠ Optimized for filter-and-aggregate queries
- View-based stream processing
  - ✓ Highly performant
  - ⚠ Requires fine-tuned and invasive integration



- Historic data in stream processing engines
  - ⚠ Read-only
  - ⚠ Optimized for filter-and-aggregate queries
- View-based stream processing
  - ✓ Highly performant
  - ⚠ Requires fine-tuned and invasive integration




→ **Current solutions are complex to implement and integrate**




# Relation-Based Integration

- Relations are entry point for all data in a DBMS




# Relation-Based Integration

- Relations are entry point for all data in a DBMS
- Database systems utilize physical data independence
  - ✓ Masks remote data access
  - ✓ Independence of storage formats, e.g., Parquet
  - ✓ Optimized access paths for OLTP or OLAP




- Relations are entry point for all data in a DBMS
- Database systems utilize physical data independence
  - ✓ Masks remote data access
  - ✓ Independence of storage formats, e.g., Parquet
  - ✓ Optimized access paths for OLTP or OLAP
- ✓ New relation types transparent to other operators




# Relation-Based Integration

- Relations are entry point for all data in a DBMS
- Database systems utilize physical data independence
  - ✓ Masks remote data access
  - ✓ Independence of storage formats, e.g., Parquet
  - ✓ Optimized access paths for OLTP or OLAP
- ✓ New relation types transparent to other operators
- ⚠ Streams are ephemeral in nature
- ⚠ Scanned data is unknown at query start




# Masking Unknown Scan Boundaries

→ Scans rely on fixed set of accessed data, i.e. TIDs




# Masking Unknown Scan Boundaries

- Scans rely on fixed set of accessed data, i.e. TIDs
- ⚠ Stream data arrives during query processing




# Masking Unknown Scan Boundaries

- Scans rely on fixed set of accessed data, i.e. TIDs
- ⚠ Stream data arrives during query processing
- Scan has to detect stream depletion
  - ⚠ Control messages would break SQL interface
  - ✓ Detect stream depletion from metadata



- Scans rely on fixed set of accessed data, i.e. TIDs
- ⚠ Stream data arrives during query processing
- Scan has to detect stream depletion
  - ⚠ Control messages would break SQL interface
  - ✓ Detect stream depletion from metadata




→ Queries with session-window semantics




# Masking Data Ephemerality

→ Scans rely on data to be fully and durably materialized





# Masking Data Ephemerality

- ➔ Scans rely on data to be fully and durably materialized
- ⚠ Streams are too voluminous to be materialized fully




# Masking Data Ephemerality

- Scans rely on data to be fully and durably materialized
- ⚠ Streams are too voluminous to be materialized fully
- Scans have to deal with short-lived data
  - ✓ Ring buffer provides required abstraction
  - ⚠ Requires careful overflow checks for variable-sized data




- Scans rely on data to be fully and durably materialized
- ⚠ Streams are too voluminous to be materialized fully
- Scans have to deal with short-lived data
  - ✓ Ring buffer provides required abstraction
  - ⚠ Requires careful overflow checks for variable-sized data




→ Ring-buffered cache with overflow handling

## Query throughput with increasing thread count



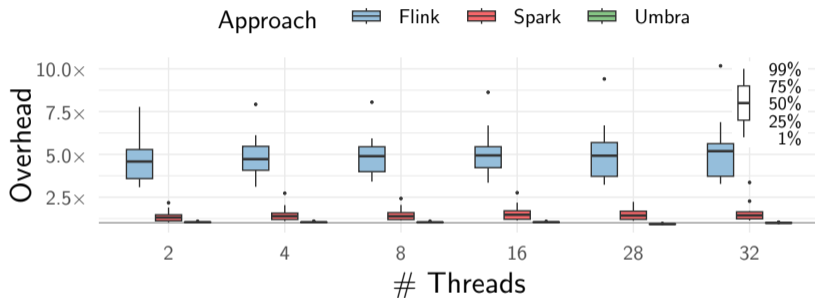
➔ TPC-H SF 100 where lineitem is treated as a stream

## Query throughput with increasing thread count




- ➔ TPC-H SF 100 where lineitem is treated as a stream
- ✓ Near-linear scaling to all physical cores
- ✓ **3×** speedup over established stream processing engines

## Runtime overhead of attaching a query over data ingestion



➔ Runtime difference of evaluating a query to parsing the data in each system


## Runtime overhead of attaching a query over data ingestion




- ➔ Runtime difference of evaluating a query to parsing the data in each system
- ✓ Next to no overhead over insert processing
- ✓ Attaching multiple queries has no additional cost




- Relation-based stream processing
  - ✓ Ring-buffered cache for ephemeral data
  - ✓ Session-windowed semantics masks unknown stream bounds
  - ✓ Minimally invasive, fully SQL-based integration
  - ✓ Full functionality for historic data



- Relation-based stream processing
  - ✓ Ring-buffered cache for ephemeral data
  - ✓ Session-windowed semantics masks unknown stream bounds
  - ✓ Minimally invasive, fully SQL-based integration
  - ✓ Full functionality for historic data
- ✓ Higher ease of integration than view-based solutions
- ✓ More performant than stream processing engines



- Relation-based stream processing
  - ✓ Ring-buffered cache for ephemeral data
  - ✓ Session-windowed semantics masks unknown stream bounds
  - ✓ Minimally invasive, fully SQL-based integration
  - ✓ Full functionality for historic data
- ✓ Higher ease of integration than view-based solutions
- ✓ More performant than stream processing engines



Thank you for your attention!