

## HDES: A Dynamic Stream Processing Engine

Nico Dulhardt, Torben Meyer, Marvin Thiele, Anton von Weltzien

14.04.2020

Masterprojekt WS 19/20  
Data Engineering Systems

# Agenda

---



1. Goals
2. Features
3. Architecture Overview
4. Query Transformation
5. Query Execution
6. Ad-hoc join processing
7. Benchmark Results

# Goals

---



1. Build a **standalone prototype** of a stream processing engine that has first class support for dynamic query deployment and removal
2. Support processing **simple queries** and **streams**
3. Support **online optimizations** for efficient multi-query processing

# Features

---



- Stream Processing Framework written in **Java 11**
- **Ad-hoc** addition and removal of arbitrary queries
- **Single** node, **multi-threaded** execution
- **Optimization** for Joins and Aggregations in **multi-query execution**
- Queries are defined in a Flink-like **dataflow language**
- Support for **Sliding-** and **Tumbling-Windows** both with **Event-** and **Processing-Time**

# Dataflow API Overview

Stream type	Name	Description	Resulting type
AStream	flatMap	Maps an incoming elements to an iterable of any type.	AStream
	map	Maps an incoming element to a arbitrary type.	AStream
	filter	Retains only the elements, which comply with a given predicate.	AStream
	window	Windows the stream with a given window assigner.	WindowedAStream
	to	Writes the elements of the stream into a sink.	None
WindowedAStream	groupBy	Groups the stream by a key.	KeyedWindowedAstream
	aggregate	Aggregates the stream in a window.	AStream
	join	Performs a join with a different stream.	AStream
	ajoin	Performs an AJoin with a different stream.	AStream
KeyedWindowedAStream	aggregate	Aggregates the stream by a key in a window.	AStream

# Code Example

```
JobManager jobManager = new JobManager();
jobManager.runEngine();
NetworkSource nws1 = new NetworkSource(7001, ...);
NetworkSource nws2 = new NetworkSource(7002, ...);

TopologyBuilder builder = TopologyBuilder.newQuery();

AStream<Tuple3<String,Float,Long>> s1 = builder.streamOf(nws1);
AStream<Tuple3<String,Float,Long>> s2 = builder.streamOf(nws2);

s1.window(TumblingWindow.ofEventTime(Time.seconds(5)))
    .join(s2,
        (t1, t2) -> new Tuple4<>(t1.v1, t1.v2, t2.v2, t1.v3),
        Tuple3::v1,
        Tuple3::v1,
        WatermarkGenerator.seconds(1, 1_000),
        t3 -> t3.v4
    )
    .to(new FileSink("join"));

Query joinQuery = builder.buildAsQuery();
jobManager.addQuery(joinQuery, 50, ChronoUnit.Seconds);
```

# Code Example

```
JobManager jobManager = new JobManager();
jobManager.runEngine();
NetworkSource nws1 = new NetworkSource(7001, ...);
NetworkSource nws2 = new NetworkSource(7002, ...);

TopologyBuilder builder = TopologyBuilder.newQuery();

AStream<Tuple3<String,Float,Long>> s1 = builder.streamOf(nws1);
AStream<Tuple3<String,Float,Long>> s2 = builder.streamOf(nws2);

s1.window(TumblingWindow.ofEventTime(Time.seconds(5)))
    .join(s2,
        (t1, t2) -> new Tuple4<>(t1.v1, t1.v2, t2.v2, t1.v3),
        Tuple3::v1,
        Tuple3::v1,
        WatermarkGenerator.seconds(1, 1_000),
        t3 -> t3.v4
    )
    .to(new FileSink("join"));

Query joinQuery = builder.buildAsQuery();
jobManager.addQuery(joinQuery, 50, ChronoUnit.Seconds);
```



Create JobManager and start engine

# Code Example

```
JobManager jobManager = new JobManager();
jobManager.runEngine();
NetworkSource nws1 = new NetworkSource(7001, ...);
NetworkSource nws2 = new NetworkSource(7002, ...);

TopologyBuilder builder = TopologyBuilder.newQuery();

AStream<Tuple3<String,Float,Long>> s1 = builder.streamOf(nws1);
AStream<Tuple3<String,Float,Long>> s2 = builder.streamOf(nws2);

s1.window(TumblingWindow.ofEventTime(Time.seconds(5)))
    .join(s2,
        (t1, t2) -> new Tuple4<>(t1.v1, t1.v2, t2.v2, t1.v3),
        Tuple3::v1,
        Tuple3::v1,
        WatermarkGenerator.seconds(1, 1_000),
        t3 -> t3.v4
    )
    .to(new FileSink("join"));

Query joinQuery = builder.buildAsQuery();
jobManager.addQuery(joinQuery, 50, ChronoUnit.Seconds);
```

A yellow rectangular box with the text "Define Sources" inside. A yellow arrow points from the box to the two NetworkSource lines in the code above.

Define Sources



# Code Example

```
JobManager jobManager = new JobManager();
jobManager.runEngine();
NetworkSource nws1 = new NetworkSource(7001, ...);
NetworkSource nws2 = new NetworkSource(7002, ...);

TopologyBuilder builder = TopologyBuilder.newQuery();

AStream<Tuple3<String,Float,Long>> s1 = builder.streamOf(nws1);
AStream<Tuple3<String,Float,Long>> s2 = builder.streamOf(nws2);

s1.window(TumblingWindow.ofEventTime(Time.seconds(5)))
    .join(s2,
        (t1, t2) -> new Tuple4<>(t1.v1, t1.v2, t2.v2, t1.v3),
        Tuple3::v1,
        Tuple3::v1,
        WatermarkGenerator.seconds(1, 1_000),
        t3 -> t3.v4
    )
    .to(new FileSink("join"));

Query joinQuery = builder.buildAsQuery();
jobManager.addQuery(joinQuery, 50, ChronoUnit.Seconds);
```

Create new query with TopologyBuilder

A yellow arrow originates from the text box and points to the line `TopologyBuilder builder = TopologyBuilder.newQuery();` in the code block above.

# Code Example

```
JobManager jobManager = new JobManager();
jobManager.runEngine();
NetworkSource nws1 = new NetworkSource(7001, ...);
NetworkSource nws2 = new NetworkSource(7002, ...);

TopologyBuilder builder = TopologyBuilder.newQuery();
```

```
    AStream<Tuple3<String,Float,Long>> s1 = builder.streamOf(nws1);
    AStream<Tuple3<String,Float,Long>> s2 = builder.streamOf(nws2);

    s1.window(TumblingWindow.ofEventTime(Time.seconds(5)))
        .join(s2,
            (t1, t2) -> new Tuple4<>(t1.v1,t1.v2,t2.v2,t1.v3),
            Tuple3::v1,
            Tuple3::v1,
            WatermarkGenerator.seconds(1, 1_000),
            t3 -> t3.v4
        )
        .to(new FileSink("join"));
```

```
Query joinQuery = builder.buildAsQuery();
jobManager.addQuery(joinQuery, 50, ChronoUnit.Seconds);
```

A yellow rectangular box with a black border containing the text "Define query". A yellow arrow points from the box to the curly braces surrounding the main code block.

Define query

# Code Example

```
JobManager jobManager = new JobManager();
jobManager.runEngine();
NetworkSource nws1 = new NetworkSource(7001, ...);
NetworkSource nws2 = new NetworkSource(7002, ...);

TopologyBuilder builder = TopologyBuilder.newQuery();

AStream<Tuple3<String,Float,Long>> s1 = builder.streamOf(nws1);
AStream<Tuple3<String,Float,Long>> s2 = builder.streamOf(nws2);

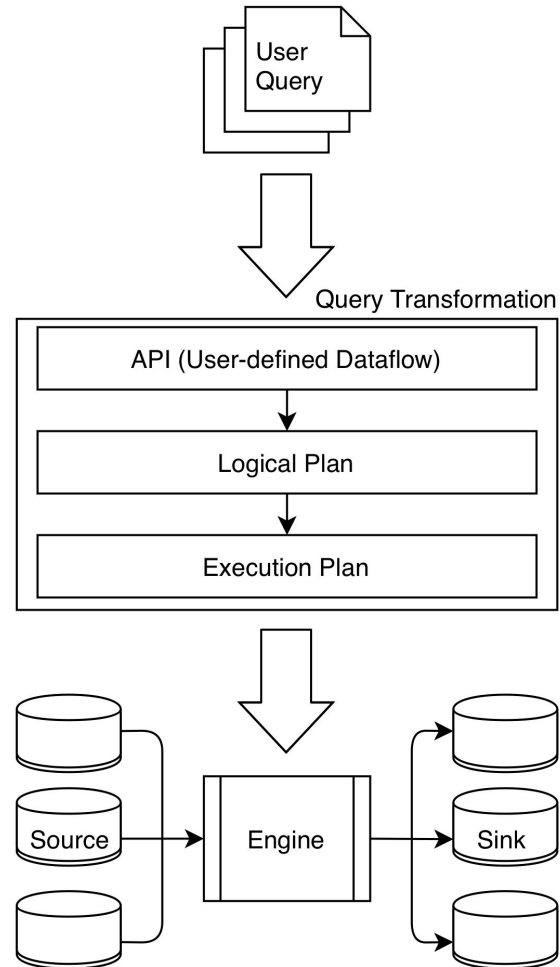
s1.window(TumblingWindow.ofEventTime(Time.seconds(5)))
    .join(s2,
        (t1, t2) -> new Tuple4<>(t1.v1, t1.v2, t2.v2, t1.v3),
        Tuple3::v1,
        Tuple3::v1,
        WatermarkGenerator.seconds(1, 1_000),
        t3 -> t3.v4
    )
    .to(new FileSink("join"));

Query joinQuery = builder.buildAsQuery();
jobManager.addQuery(joinQuery, 50, ChronoUnit.Seconds);
```



Build and submit query

## 3. Architecture Overview



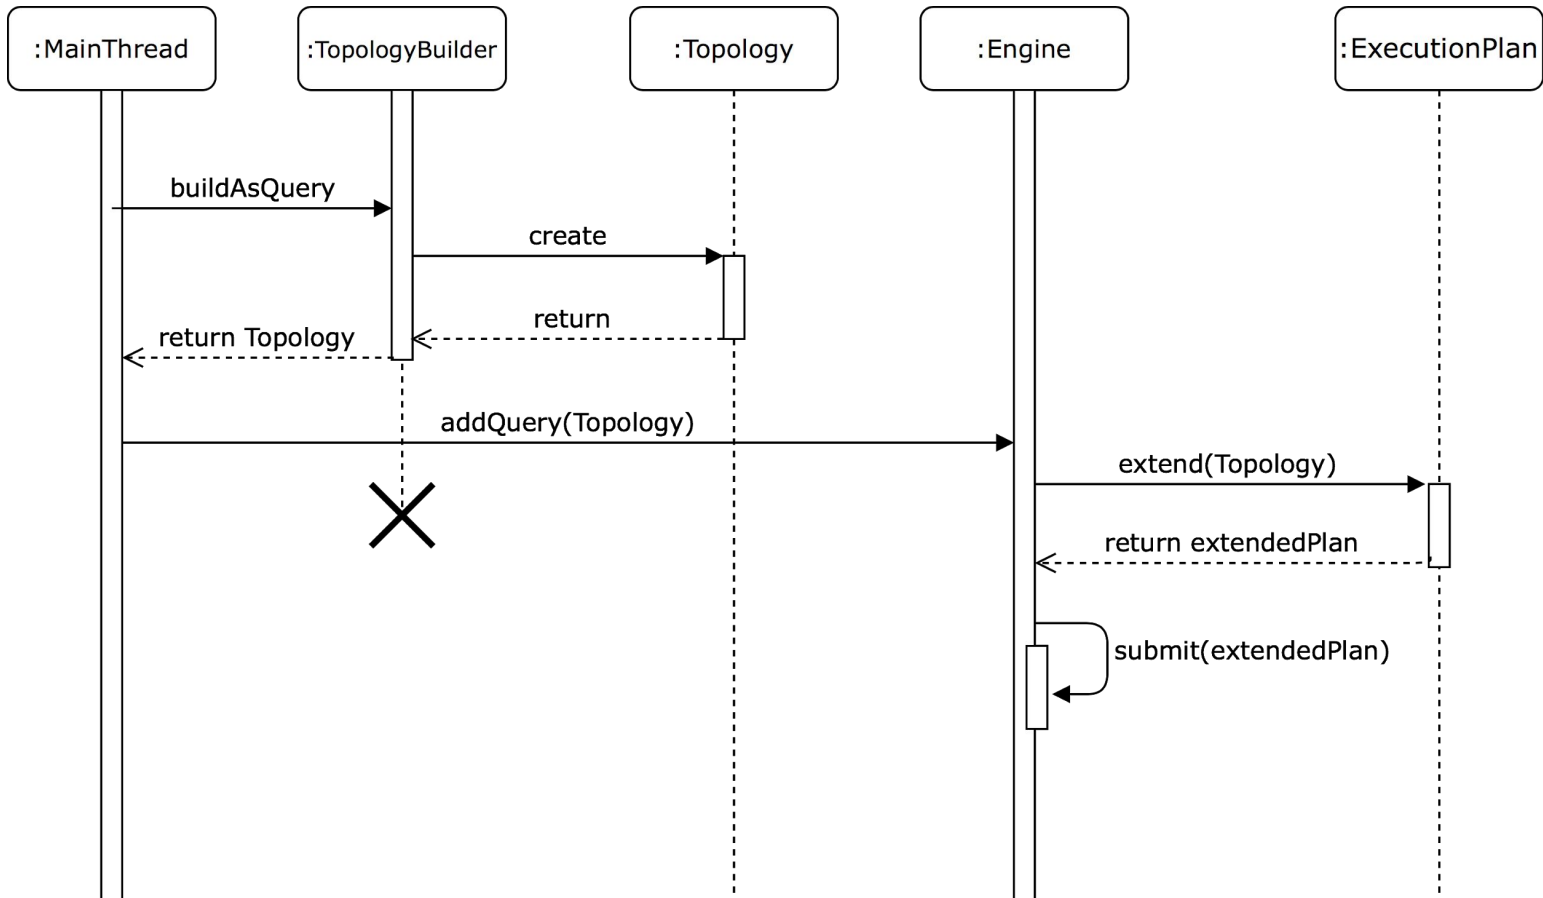
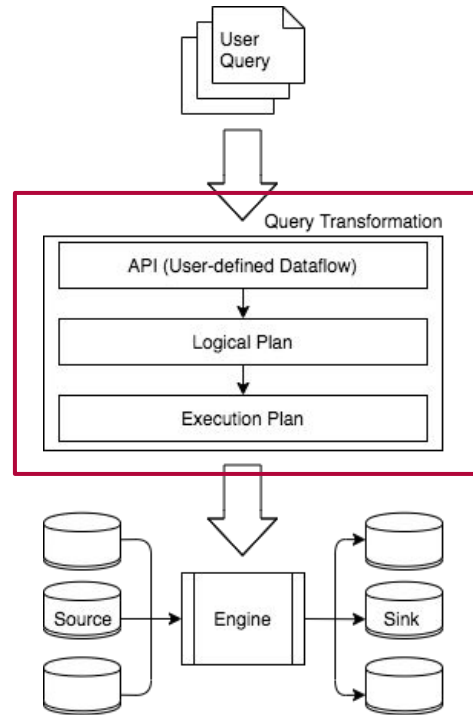


Chart 14

## 4. Query Transformation

# Transformation Pipeline





# Operators

---



## **Source:**

- Read from a source
- Attaches metadata

## **OneInputOperator:**

- Transform a single event into  $n$  new events

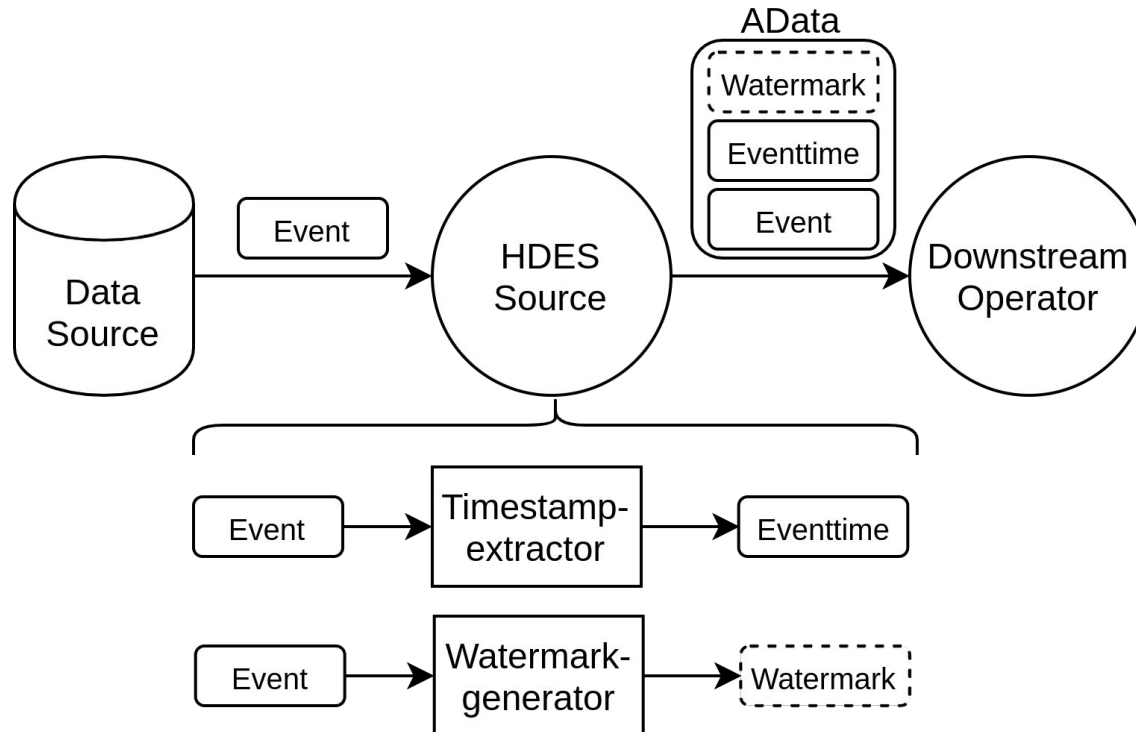
## **TwoInputOperator:**

- Transform events from two different origins

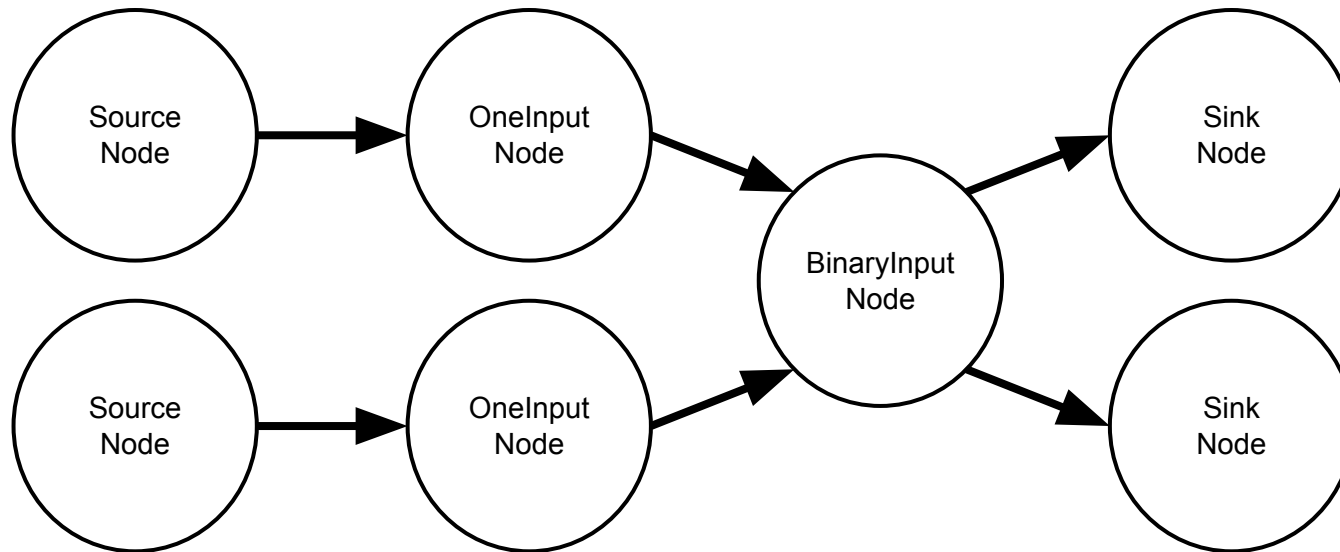
## **Sink:**

- Write to a sink

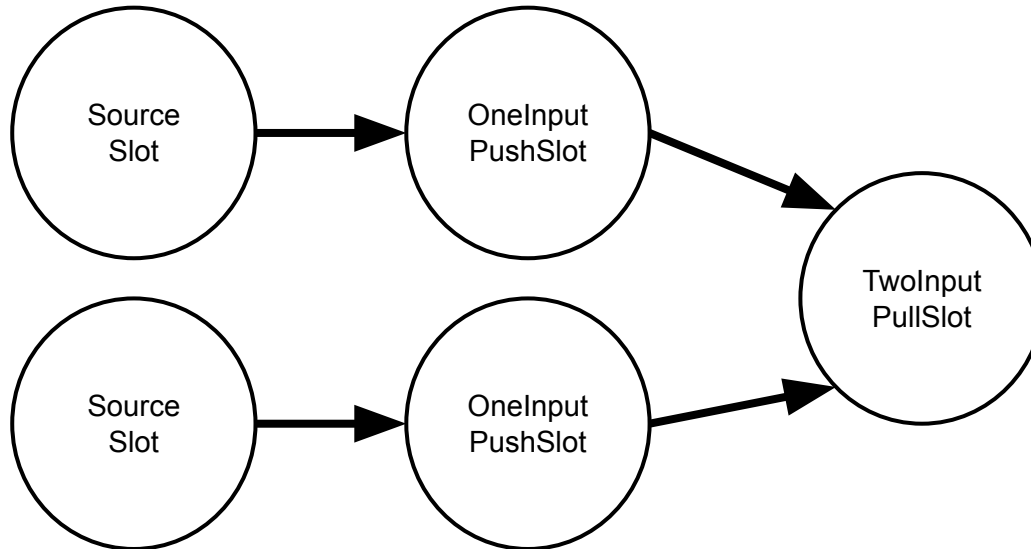
# Source Operator



# Logical Plan



# Execution Plan



# Transformation Properties

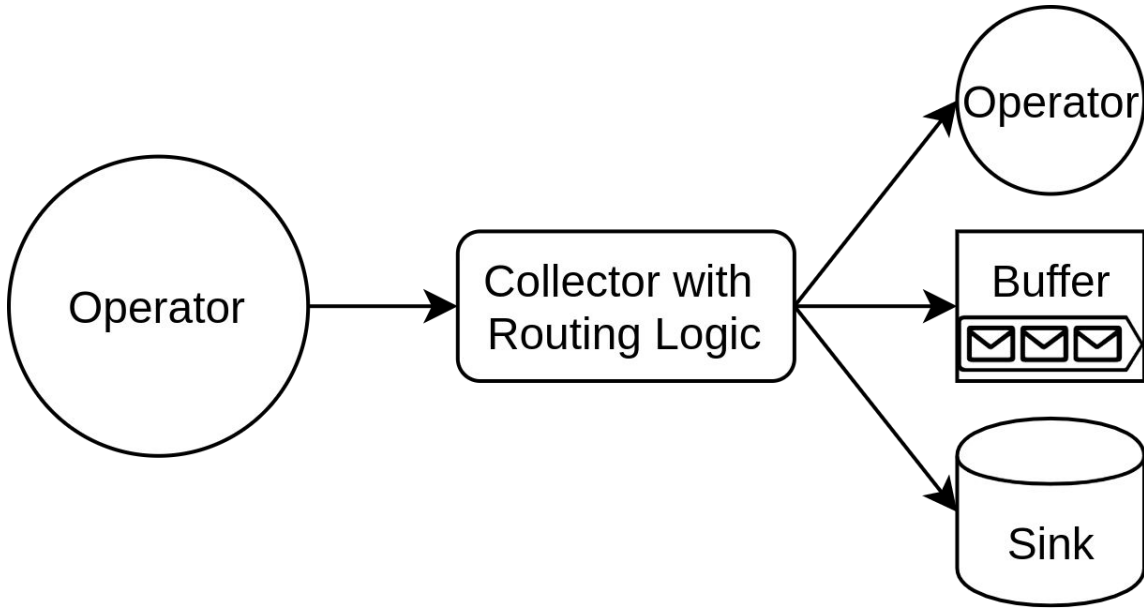
---



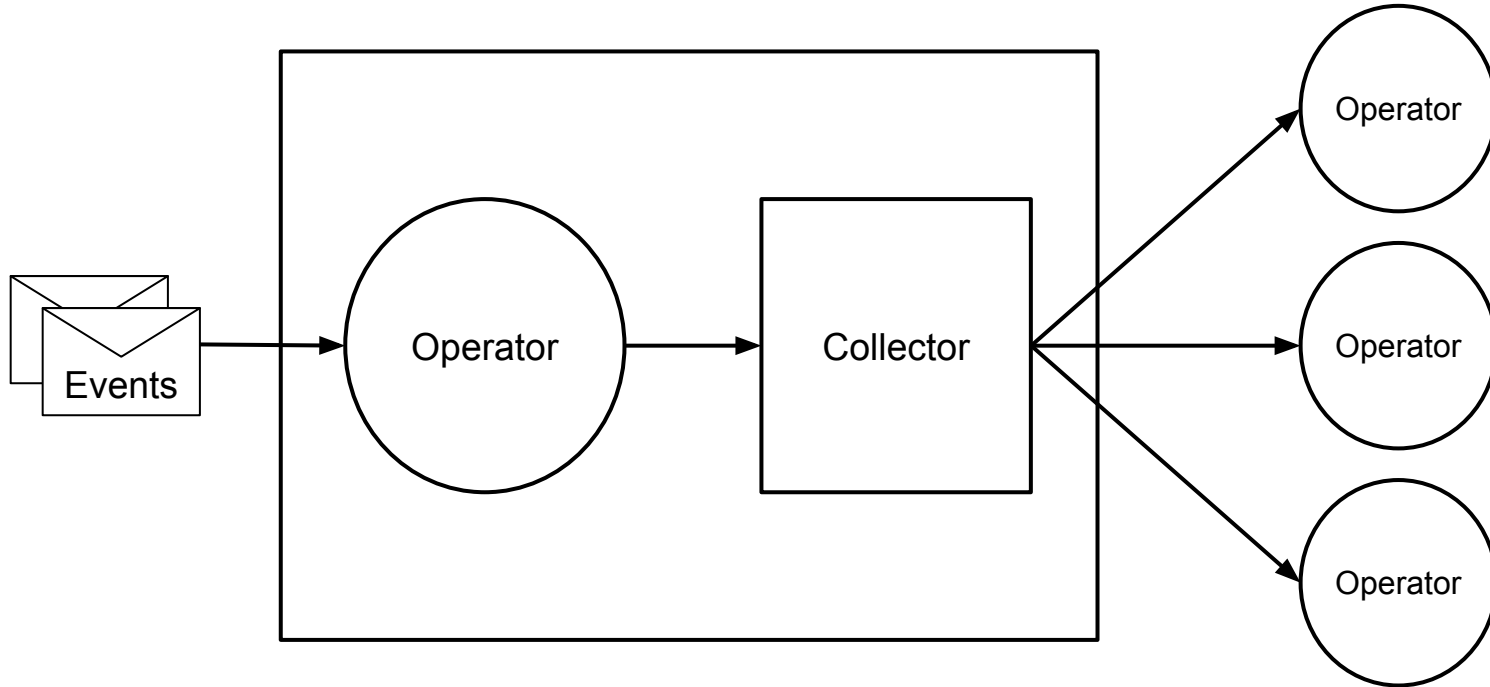
Layered architecture decouples query definition and execution.

- Interchangeable query definition
- Interchangeable Execution Plan

## 5. Query Execution




# Slot



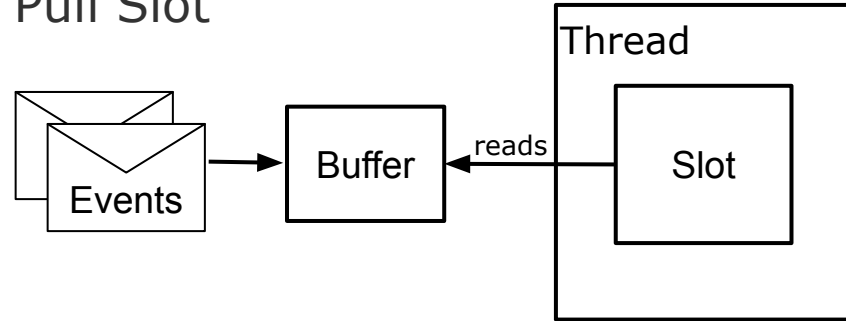


# Slot Types

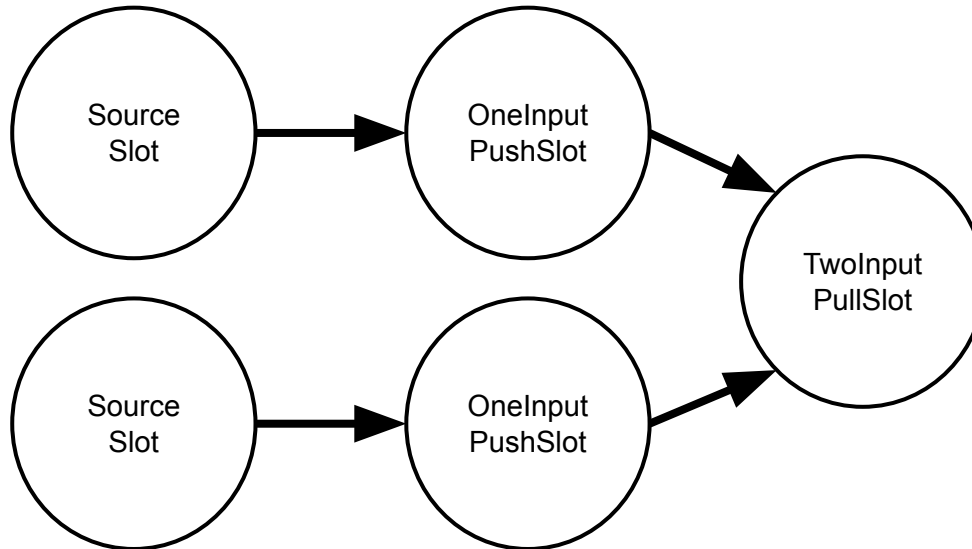
Push Slot



Pull Slot




# Execution Plan




## 6. Ad-hoc join processing

# Efficient Distributed Join Architecture




# Efficient Distributed Join Architecture




- Indexing
- Windowing

# Efficient Distributed Join Architecture



- Indexing
- Windowing
- Set intersection of join index

# Efficient Distributed Join Architecture




- Indexing
- Windowing

- Set intersection of join index

- Joins matching tuples
- Pushes to output channels

# AJoin in HDES

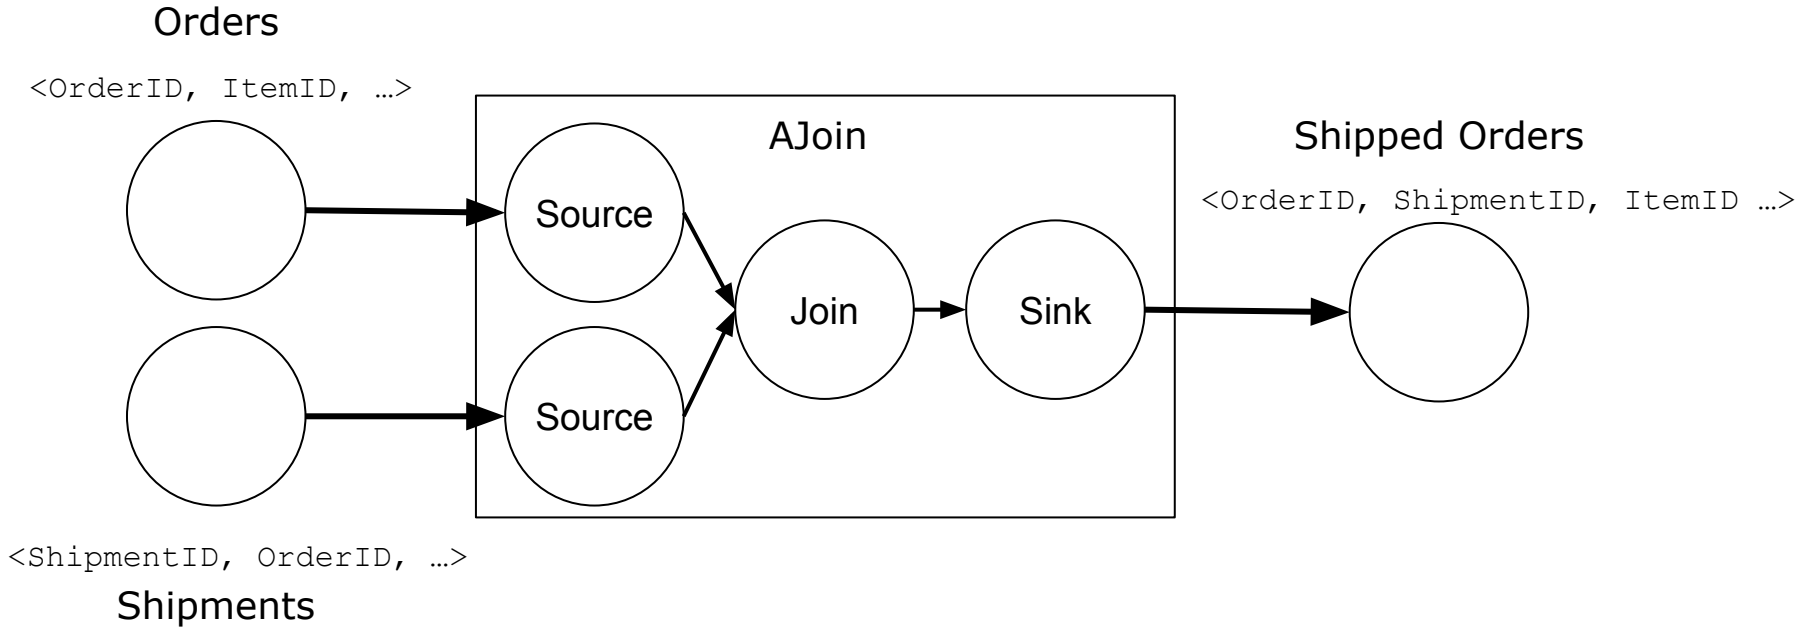
Upstream Operator





Upstream Operator




# HDES AJoin Example



# HDES AJoin Example



# HDES AJoin Example



$\langle 1, 5, \dots \rangle$     $\langle 1, 8, \dots \rangle$     $\langle 4, 7, \dots \rangle$     $\langle 5, 2, \dots \rangle$     $\langle 5, 7, \dots \rangle$

$\langle 5, 7, \dots \rangle$     $\langle 6, 8, \dots \rangle$     $\langle 9, 1, \dots \rangle$     $\langle 6, 4, \dots \rangle$     $\langle 3, 5, \dots \rangle$


1  $\leftarrow$  [ $\langle 1, 5, \dots \rangle$ ,  $\langle 1, 8, \dots \rangle$ ]  
 4  $\leftarrow$  [ $\langle 4, 7, \dots \rangle$ ]  
 5  $\leftarrow$  [ $\langle 5, 2, \dots \rangle$ ,  $\langle 5, 7, \dots \rangle$ ]

Orders Bucket

7  $\leftarrow$  [ $\langle 5, 7, \dots \rangle$ ]  
 8  $\leftarrow$  [ $\langle 6, 8, \dots \rangle$ ]  
 1  $\leftarrow$  [ $\langle 9, 1, \dots \rangle$ ]  
 4  $\leftarrow$  [ $\langle 6, 4, \dots \rangle$ ]  
 5  $\leftarrow$  [ $\langle 3, 5, \dots \rangle$ ]




Shipment Bucket

# HDES AJoin Example





1  $\leftarrow$  [ $\langle 1, 5, \dots \rangle, \langle 1, 8, \dots \rangle$ ] 4  $\leftarrow$  [ $\langle 4, 7, \dots \rangle$ ] 5  $\leftarrow$  [ $\langle 5, 2, \dots \rangle, \langle 5, 7, \dots \rangle$ ]

7  $\leftarrow$  [ $\langle 5, 7, \dots \rangle$ ] 8  $\leftarrow$  [ $\langle 6, 8, \dots \rangle$ ] 1  $\leftarrow$  [ $\langle 9, 1, \dots \rangle$ ] 4  $\leftarrow$  [ $\langle 6, 4, \dots \rangle$ ] 5  $\leftarrow$  [ $\langle 3, 5, \dots \rangle$ ]




# HDES AJoin Example



## 7. Benchmarking Results

# Benchmarking Setup

## Data Generator




16 GB RAM  
4-core Intel i5 7300HQ  
(2.50 GHz)

1 GBit/s LAN

Custom Serialization

## Engine



16 GB RAM  
4-core Intel i7 2600K  
(3.50 GHz)

# Description of Metrics

---

## Timestamps:

- Event Time = The moment the tuple is **created**
- Processing Time = The moment the tuple **enters** the **engine**
- Ejection Time = The moment the tuple **enters** the **sink**

## Latency:

- Event Time Latency = Ejection Time - Event Time
- Processing Time Latency = Ejection Time - Processing Time

## Other metrics:

- Throughput = Amount of tuples send from the data-generator per second



# Data Sets - "Basic" and "Nexmark-Light"

Legend:  
Long  
Integer

**Tuple** (12 Bytes)  
(**EventTime**, **Number**)

## Basic

Simplistic test-data

**Auction** (28 Bytes)  
(**AuctionId**, **Quantity**, **Type**, **Reserve**,  
**EventTime**)

**Bid** (32 Bytes)  
(**BidId**, **AuctionId**, **BetterId**, **Price**,  
**EventTime**)

## Nexmark-Light

More realistic data an auction  
house use-case

# Benchmark Dimensions

Test Scenario	Dataset	Query	Workload	Engine
X static-queries	Basic Nexmark	Map	1	HDES Flink
Add & Remove X-queries every 10 seconds		Filter	10	
Add a query every 10 seconds		Join	100	
		HotCat	...	
Remove a query every 10 seconds		HPPA		

# Benchmark Dimensions

Test Scenario	Dataset	Query	Workload	Engine
X static-queries	Basic	Map	1	HDES
Add & Remove X-queries every 10 seconds		Filter	10	
Add a query every 10 seconds		Join	100	
Remove a query every 10 seconds		HotCat	...	
	Nexmark	HotCat		Flink
		HPPA		

Chart 43

**Scenario: 10 fixed map-queries executed with HDES and the basic dataset**

# Static Query Execution

# Static X-Queries - HDES Map vs. Flink Map

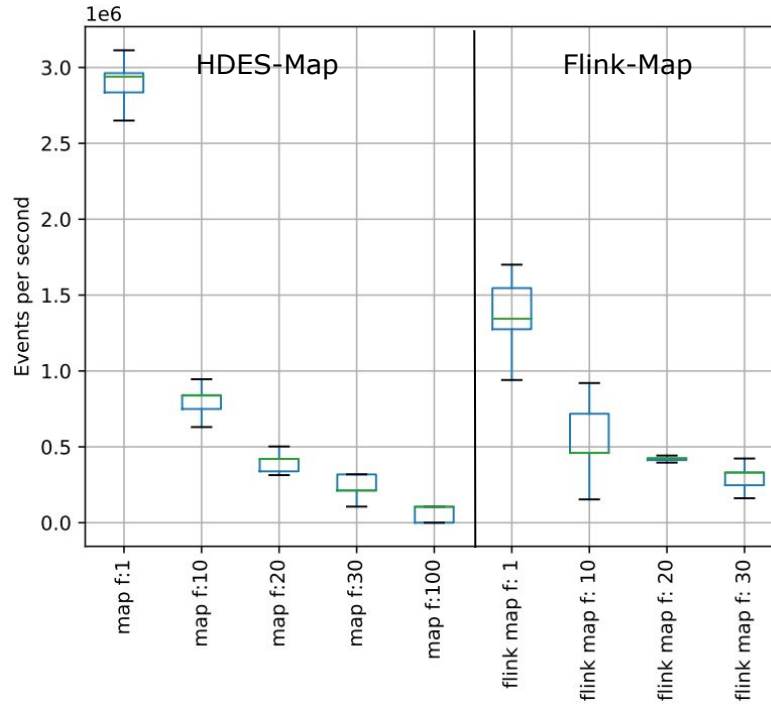
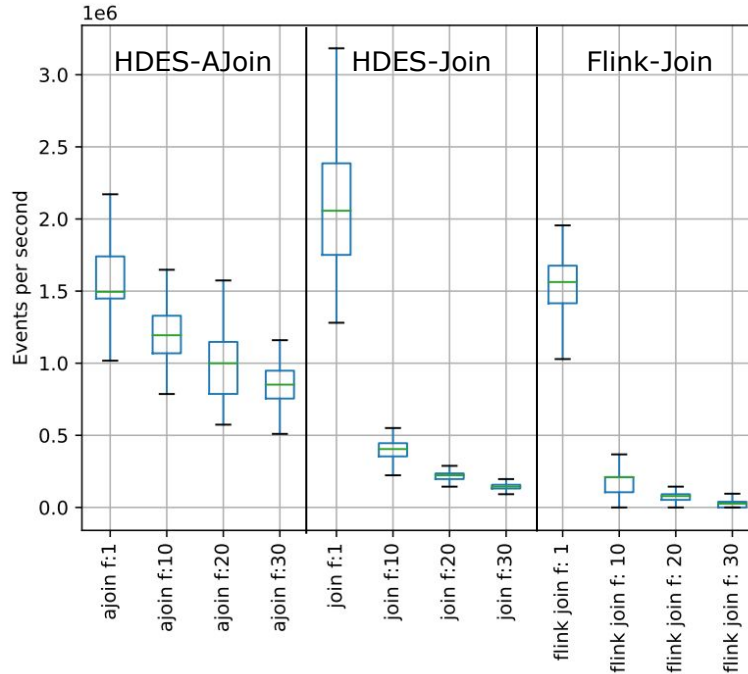


Chart 45

HDES can outperform Flink with simple operations across datasets


# Static X-Queries - HDES (A)Join vs. Flink Join



AJoin is better than Join in multi-query situations and HDES outperforms Flink's Join in every scenario

# Dynamic Query Execution

# Dynamic - Increased Adding & Removing

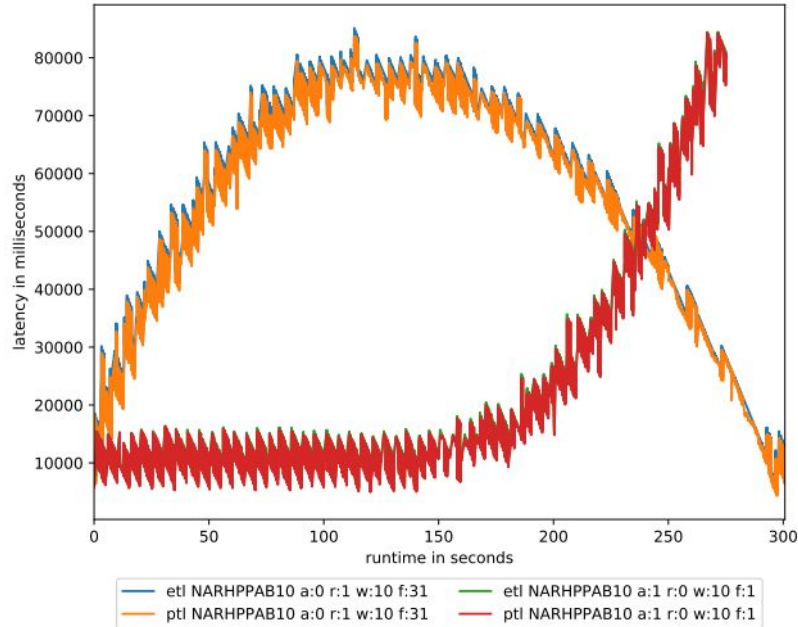


Basic - AJoin & Join

Chart 48




# Dynamic - Increased Adding & Removing



Nexmark - HPPA

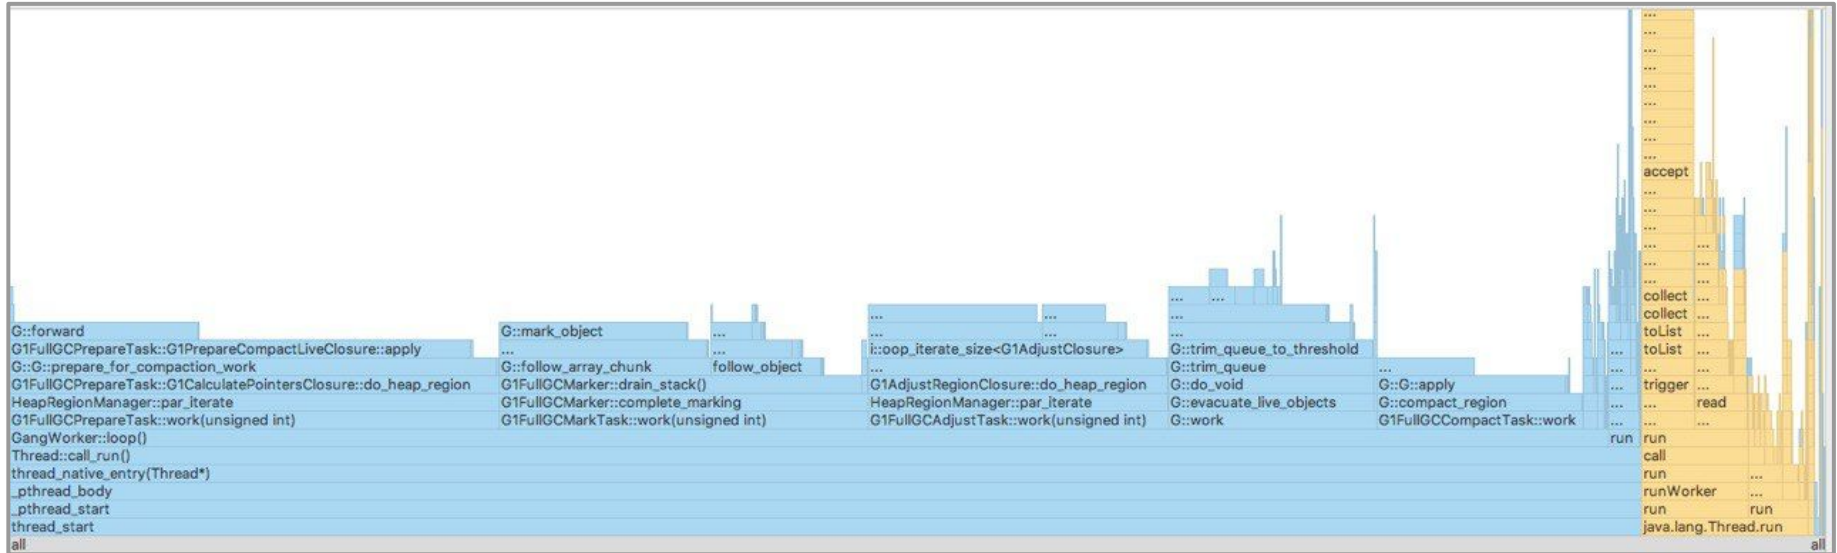
Chart 49

# Dynamic - Constant Adding & Removing



Nexmark - Hottest Category


# Garbage Collection




**Garbage Collection**

**Engine**


# Conclusion



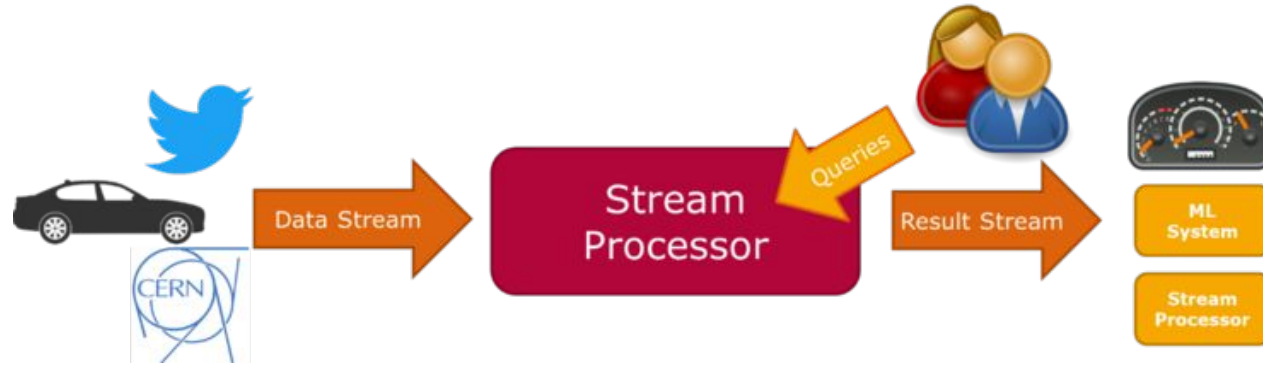
HDES **outperforms** Flink in most situations and can even compute **larger workloads**



HDES offers **ad-hoc functionality** as a first-class citizen



HDES offers integration of **optimized sharing techniques** from previous work



Thank you for your attention!

Nico Dulhardt, Torben Meyer, Marvin Thiele, Anton von Weltzien  
Masterprojekt WS 19/20  
Data Engineering Systems

# Agenda

---



1. Goals (Anton) ~ 1
2. Features (Anton) ~ 2-4
3. Architecture Overview (Anton) 2-3
4. Query Transformation (Nico) 4
5. Query Execution (Nico) 3
6. Multi-Query Execution (Torben) > 7 min
7. Benchmark Results (Marvin) >7 min

# Dynamic - Scalability of Join & AJoin


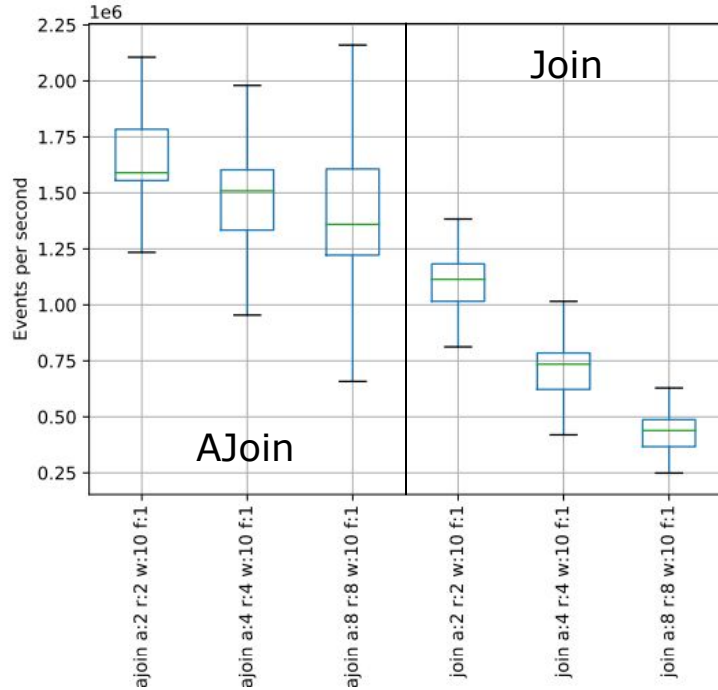


Chart 55