The Dataflow Model

A Practical Approach to Balancing Correctness, Latency, and Cost in Massive-Scale, Unbounded, Out-of-Order Data Processing

Tyler Akidau et al.

Christopher Little

Outline

Prerequisites
Problem
System
Evaluation

Prerequisites

Event vs Processing Time

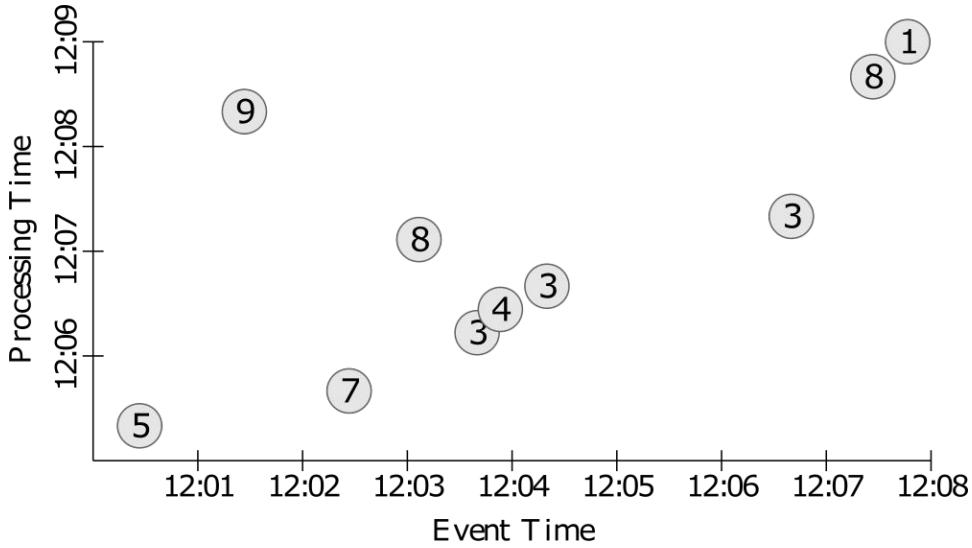

Low Watermark

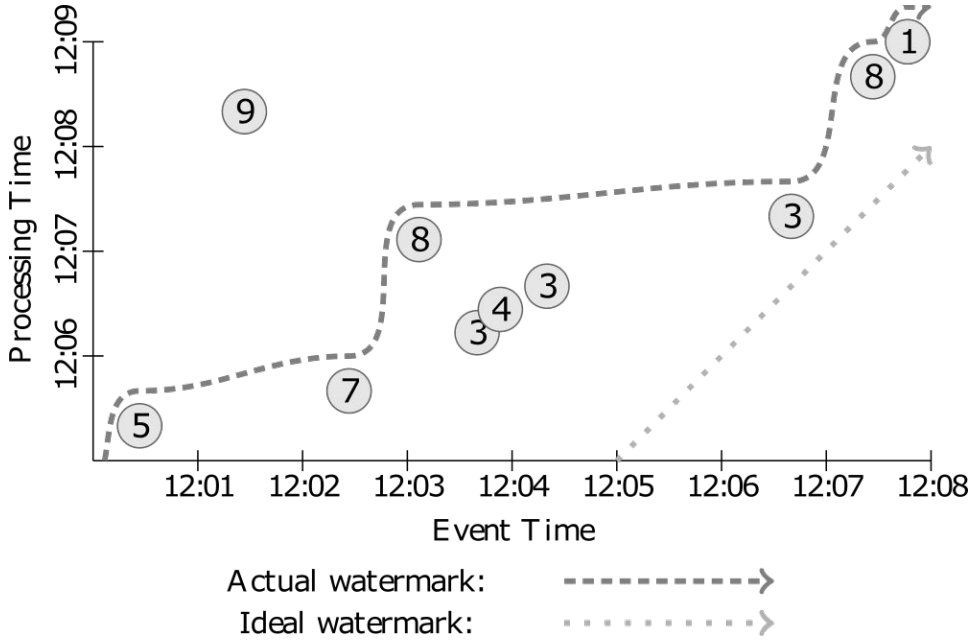

Fixed Windowing

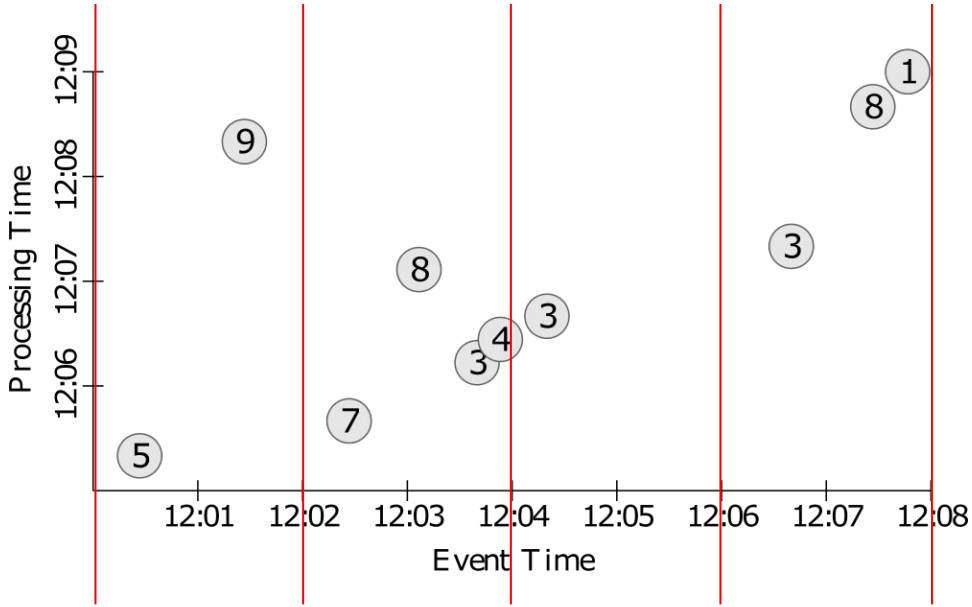

Unaligned Windowing (Tuples)

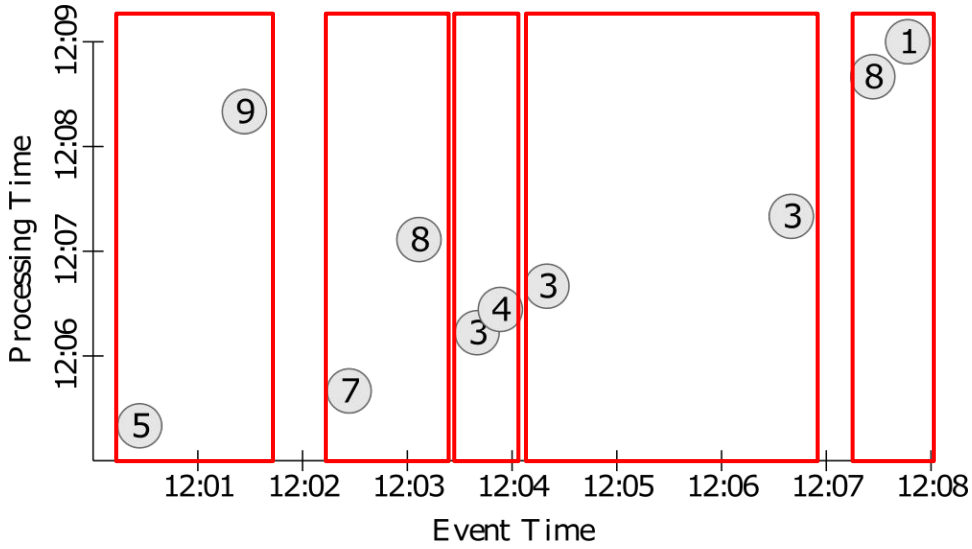

Unaligned Windowing (Sessions)

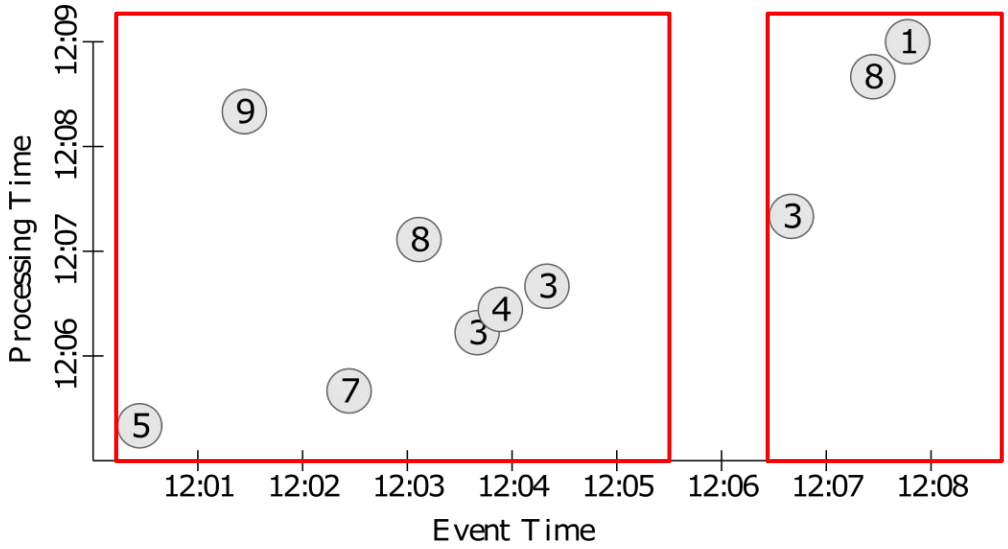

Problem

Tracking Video Sessions

- Online/**Offline** video platform
- Want aggregate stats per user: track **sessions**
- Pay advertisers per view: must be **correct**
- Want to adjust bids fast: **low latency**
- Must scale: **distributed system**

“A major shortcoming of all the models and systems mentioned above, is that they focus on input data as something which will at some point become complete.”

System

– ***What*** results are being computed.

– ***Where*** in event time they are being computed.

– ***When*** in processing time they are materialized.

– ***How*** earlier results relate to later refinements.

– ***What*** results are being computed. ✔

– ***Where*** in event time they are being computed.

– ***When*** in processing time they are materialized.

– ***How*** earlier results relate to later refinements.

Two Primitive Transforms

`(fix, 1)` `(fit, 2)`

ParDo *(ExpandPrefixes)*

`(f, 1)` `(fi, 1)` `(fix, 1)` `(f, 2)` `(fi, 2)` `(fit, 2)`

GroupByKey

`(f, [1, 2])` `(fi, [1, 2])` `(fix, [1])` `(fit, [2])`

Session Windowing Example

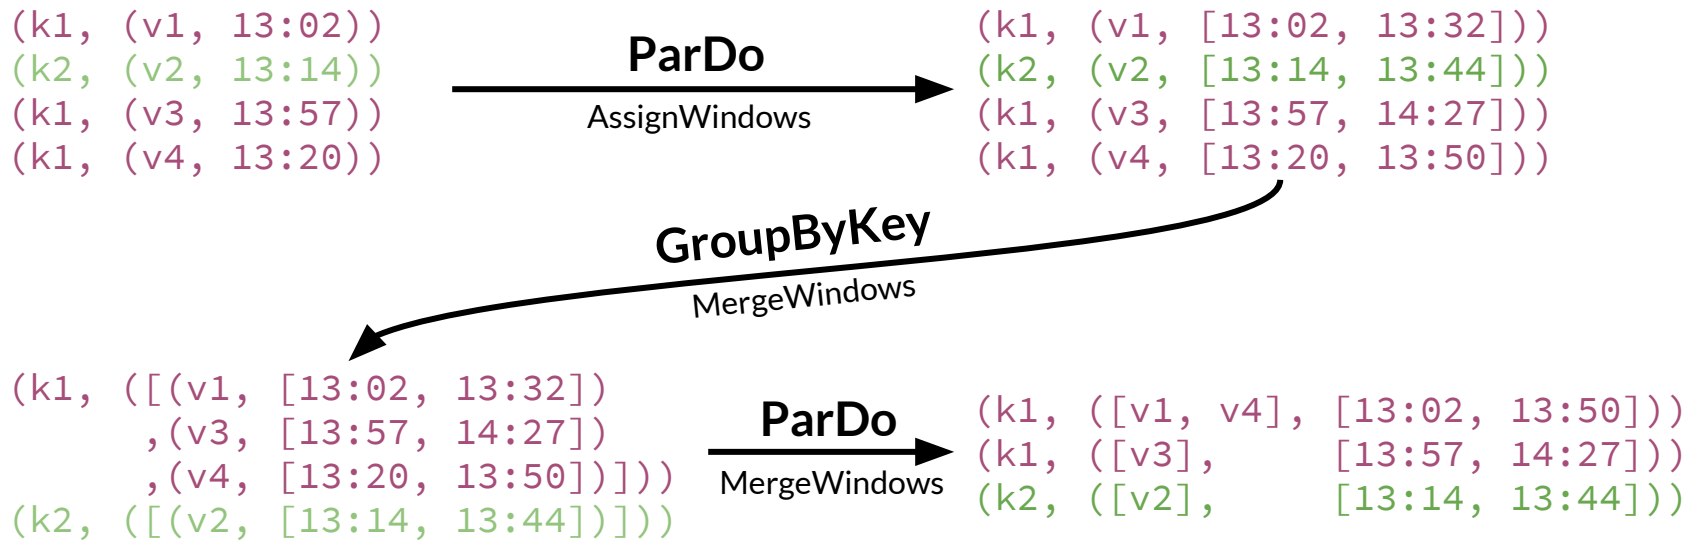

– ***What*** results are being computed. ✔

– ***Where*** in event time they are being computed. ✔

– ***When*** in processing time they are materialized.

– ***How*** earlier results relate to later refinements.

Triggering

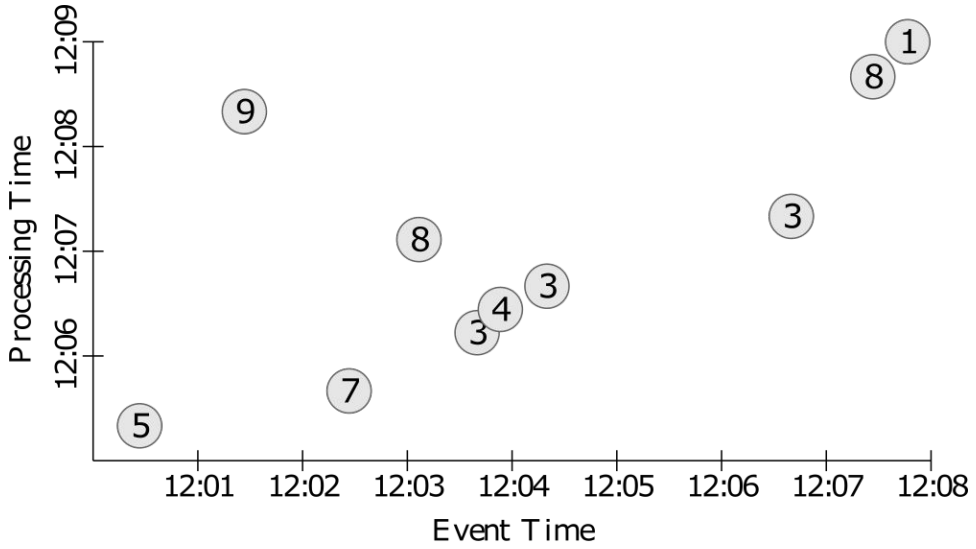
Processing Time
12:06
12:07
12:08
12:09
Event Time
12:01
12:02
12:03
12:04
12:05
12:06
12:07
12:08
5
7
9
8
3
4
3
3
8
1

Triggering (end of time)

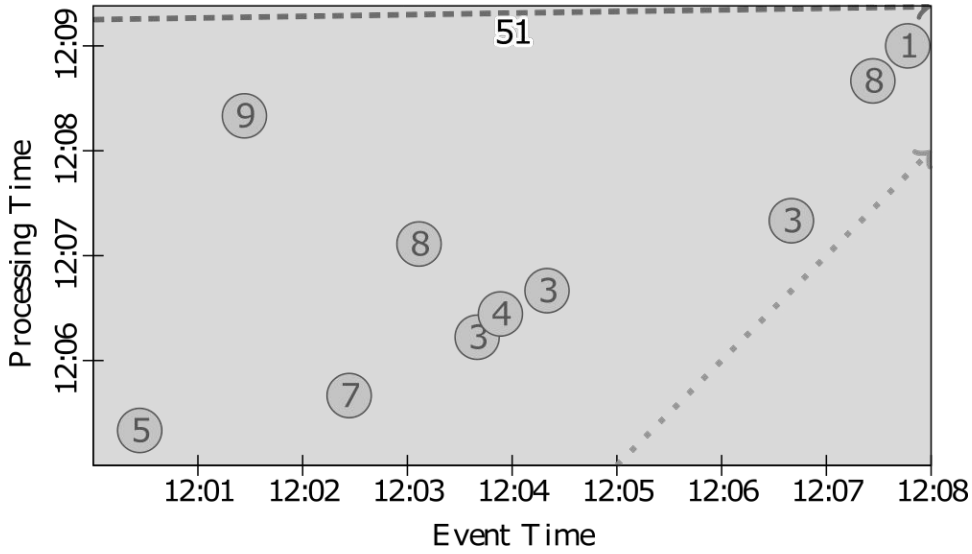

Triggering (periodically)

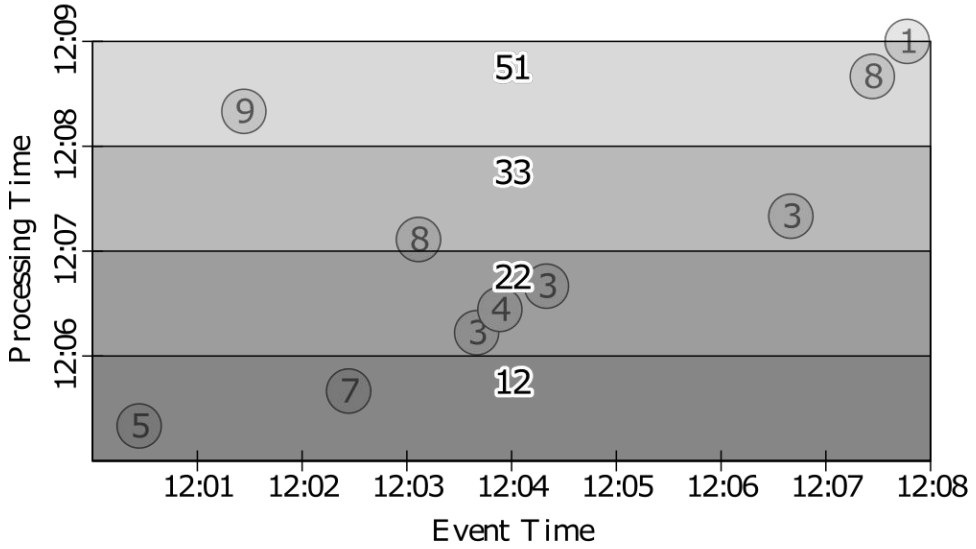

Triggering (on input, tuples)

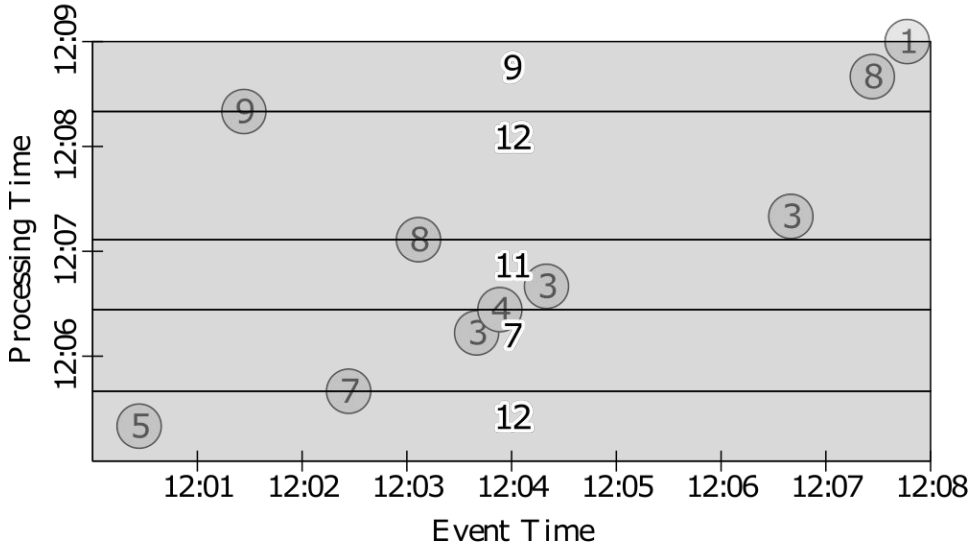

Triggering (on watermark+input)

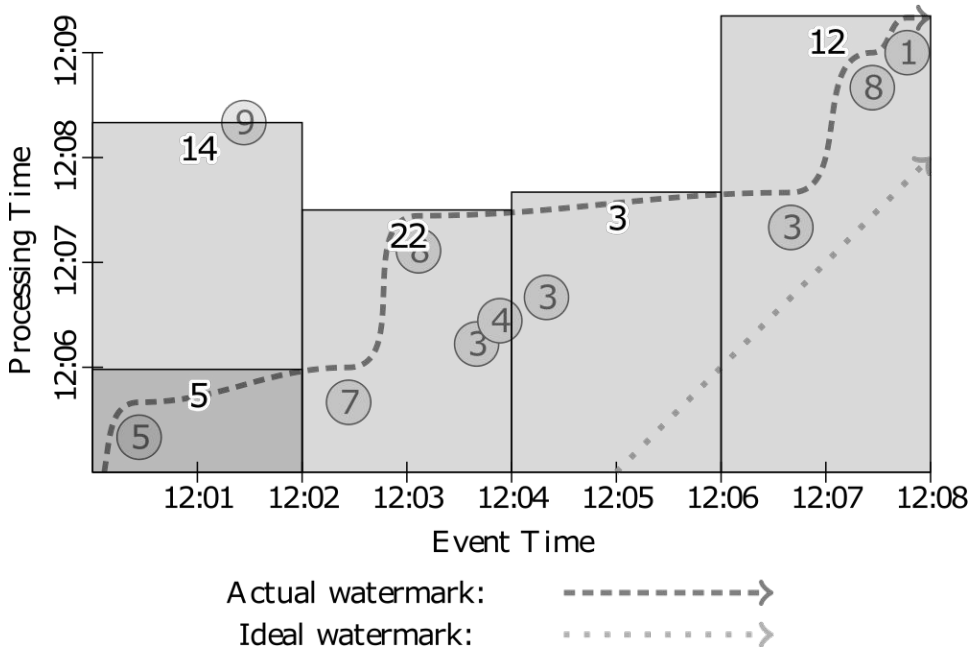

– ***What*** results are being computed. ✔

– ***Where*** in event time they are being computed. ✔

– ***When*** in processing time they are materialized. ✔

– ***How*** earlier results relate to later refinements.

Accumulating

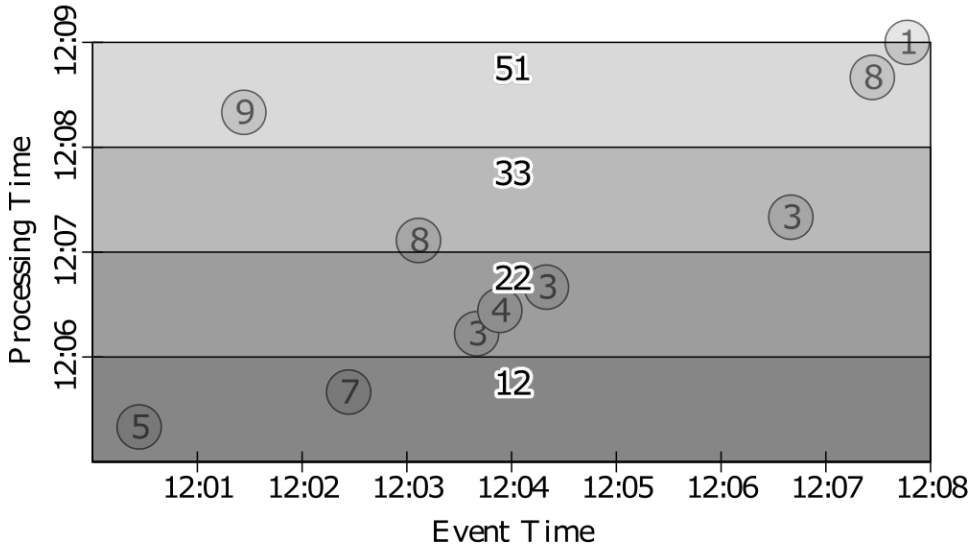

Discarding

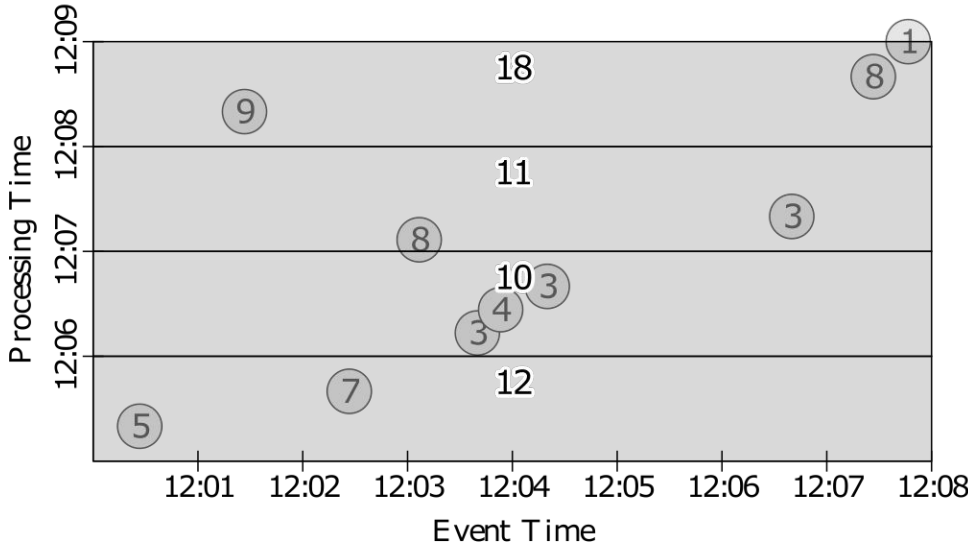
Processing Time
12:09
12:08
12:07
12:06
18
11
10
12
1
8
9
3
8
3
4
3
7
5
12:01
12:02
12:03
12:04
12:05
12:06
12:07
12:08
Event Time

Accumulating + Retracting

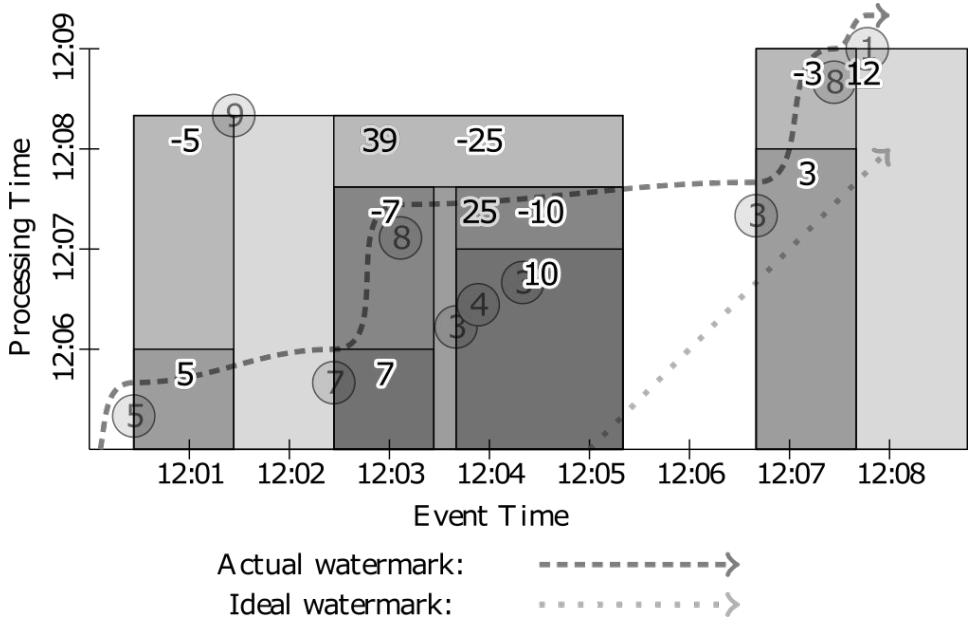

– ***What*** results are being computed. ✔

– ***Where*** in event time they are being computed. ✔

– ***When*** in processing time they are materialized. ✔

– ***How*** earlier results relate to later refinements. ✔

Evaluation

Evaluation

- Name
- Concepts
- Necessity
- Clarity

The Dataflow Model

A Practical Approach to Balancing Correctness, Latency, and Cost in Massive-Scale, Unbounded, Out-of-Order Data Processing

Tyler Akidau et al.

Christopher Little