Disco: Efficient Distributed Window Aggregation

23rd EDBT | March 30 - April 2, 2020 | Copenhagen, Denmark

Lawrence Benson, Phillip M. Grulich, Steffen Zeuch, Volker Markl, Tilmann Rabl

Contribution

Disco is an approach to distribute the aggregation of complex window types and aggregation functions.

Our main contributions are:

Disco ...

- » ... **avoids central** raw data **collection** whenever possible.
- » ... contains different window **merging** strategies to **distribute** common **window types**.
- » ... **scales linearly** and **reduces** the **network cost** significantly compared to state-of-the-art.

Current Stream Processing

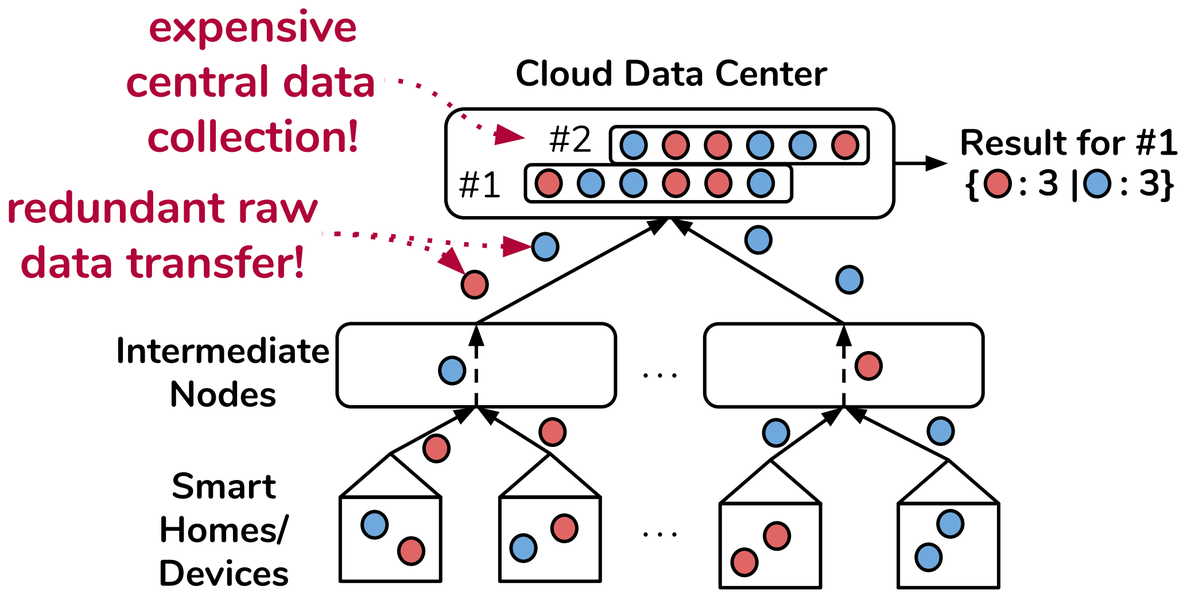

Architecture

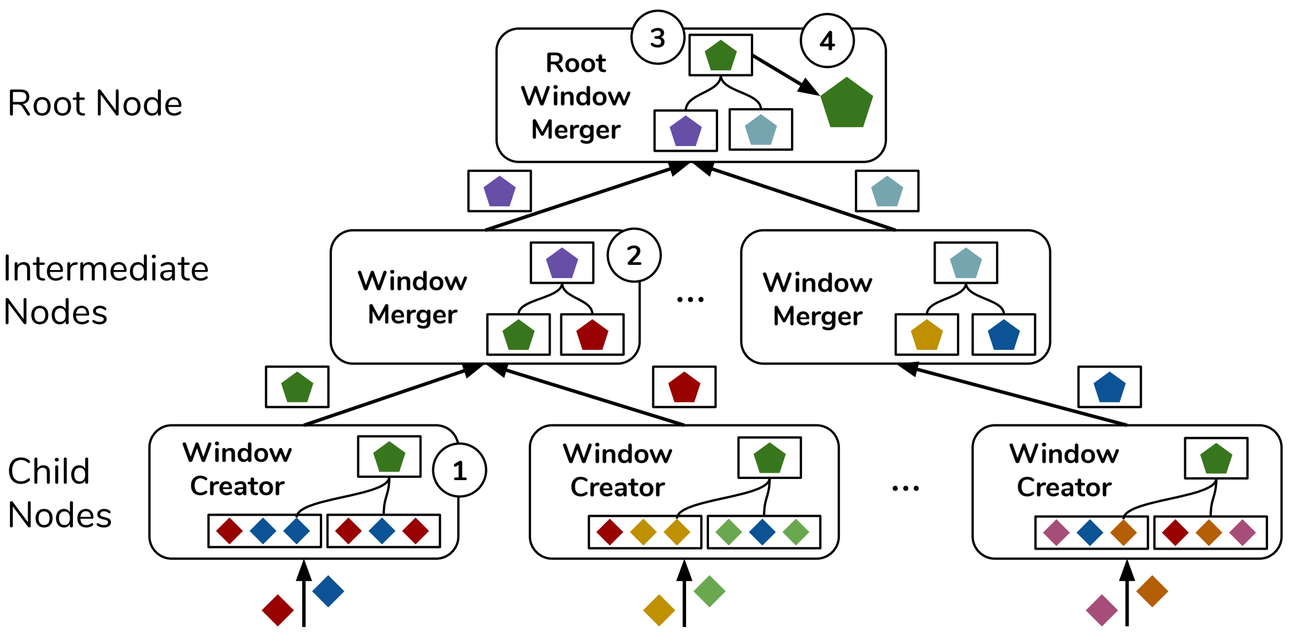

1. **Window Creator** creates slices and windows according to user-specified queries.
2. **Window Merger** merges incoming windows according to Disco's strategies.
3. **Root Window Merger** performs final window merge.
4. **Root Window Merger** returns final window aggregate.

Merging Strategies

No merging for **unique windows**.

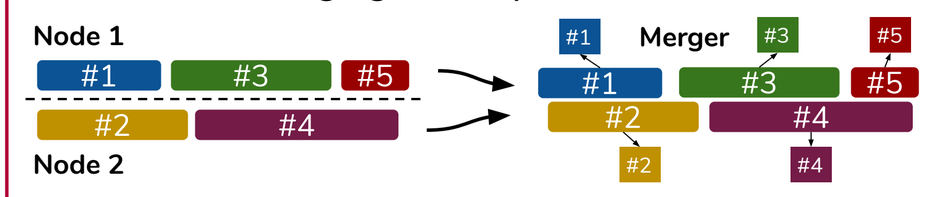

Merge **context-free windows** based on **identical bounds**.

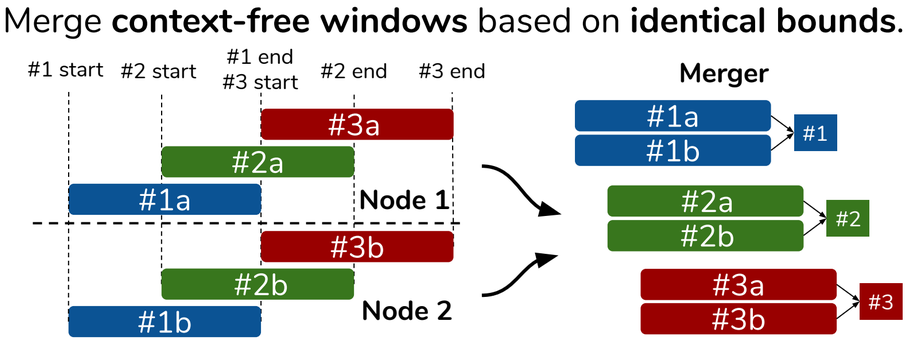

Merge **session windows** based on **global gaps**.

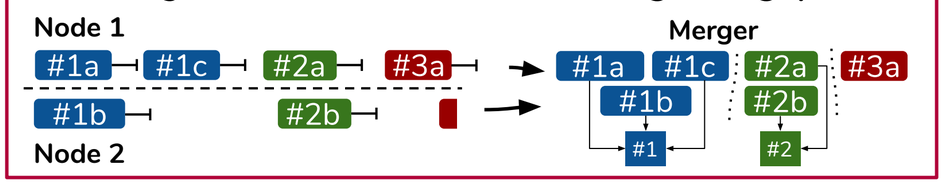

Aggregation Handling

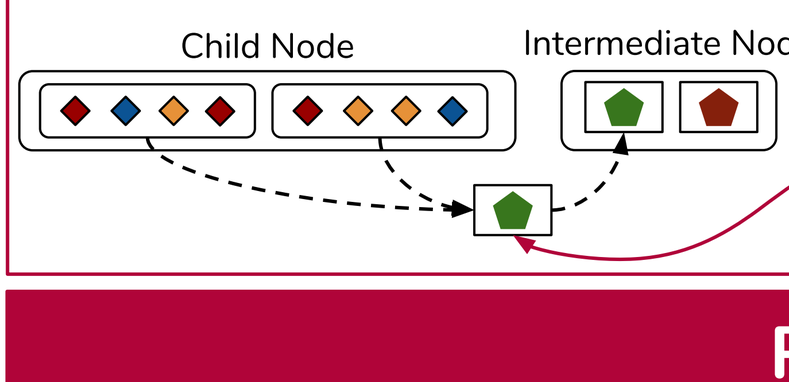

Decomposable Aggregation

transfer partial aggregates instead of raw events

Holistic Aggregation

transfer slices as smallest unit instead of individual events

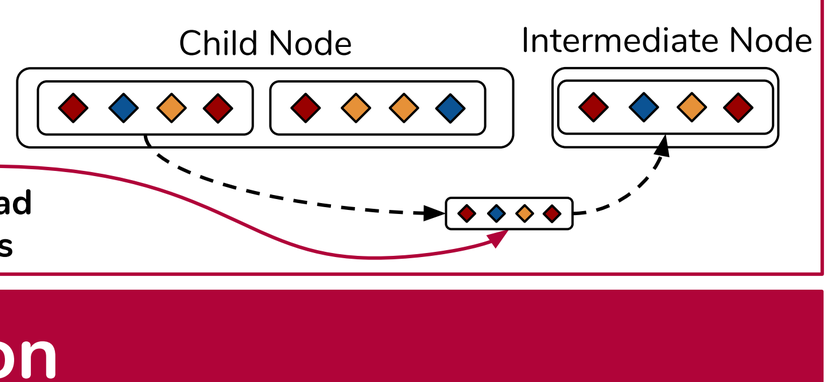

Performance Evaluation

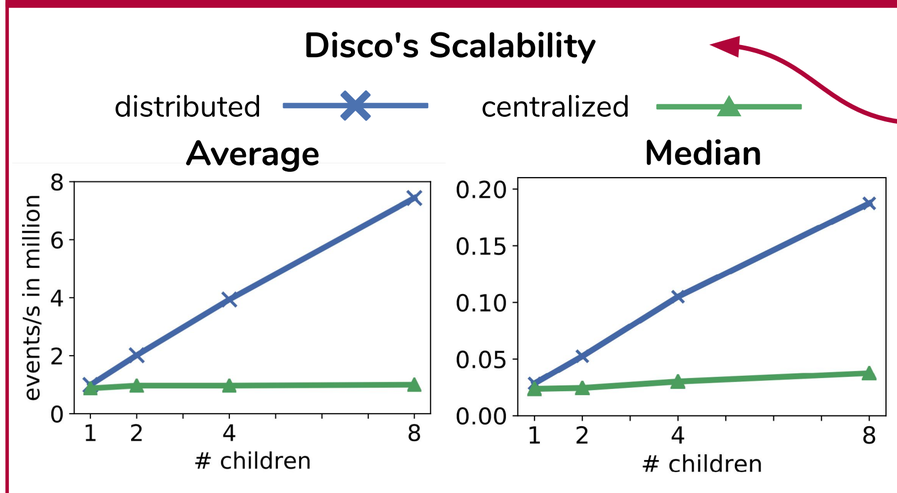

Key Results

linear scale up for both algebraic and holistic aggregation functions

significant reduction in network cost due to removed redundant raw data transfer

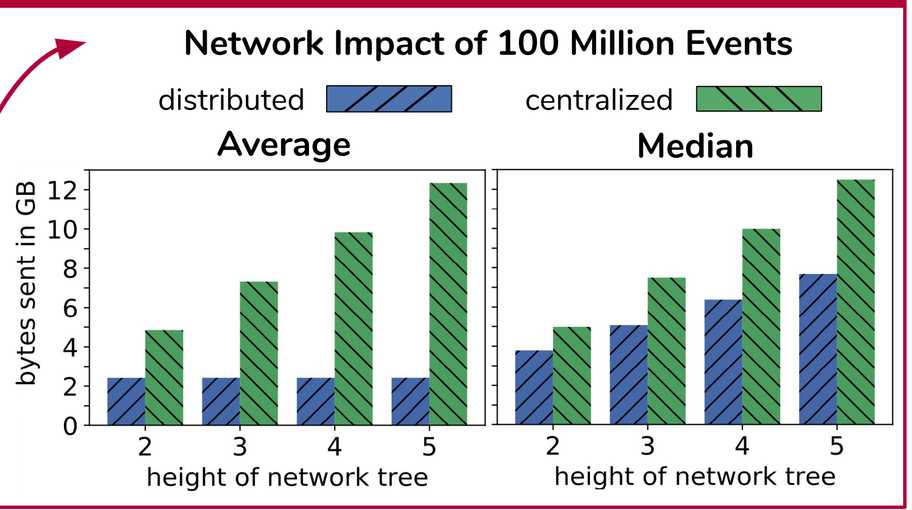

Data Engineering Systems Group
Lawrence Benson
Hasso Plattner Institute, Potsdam, Germany
E-Mail: lawrence.benson@hpi.de