Optimal and General Out-of-Order Sliding Window Aggregation

Kanat Tangwongsan[#], Martin Hirzel[+], **Scott Schneider**[+]

[#]*Mahidol University International College*
[+]*IBM T. J. Watson Research Center*

Windows on data streams

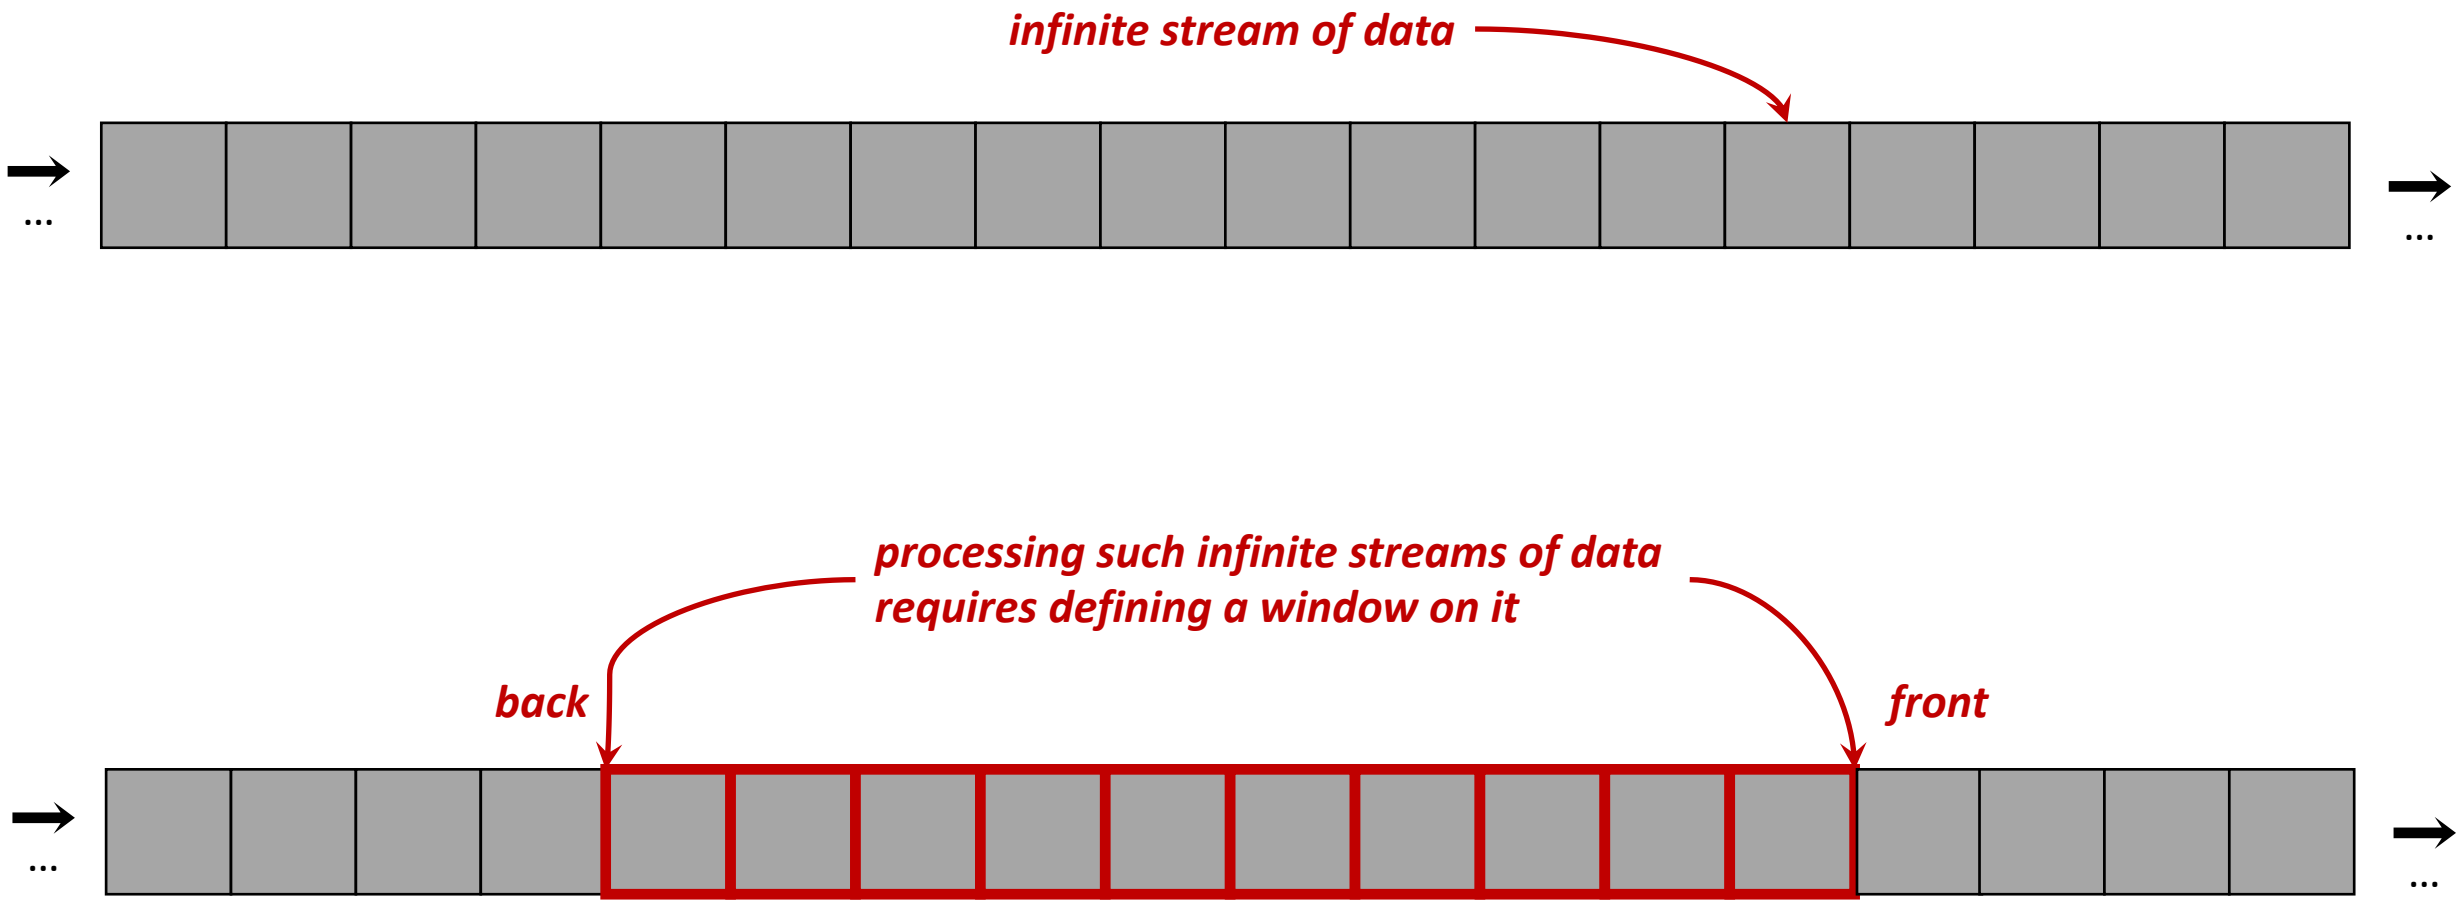

Sliding windows

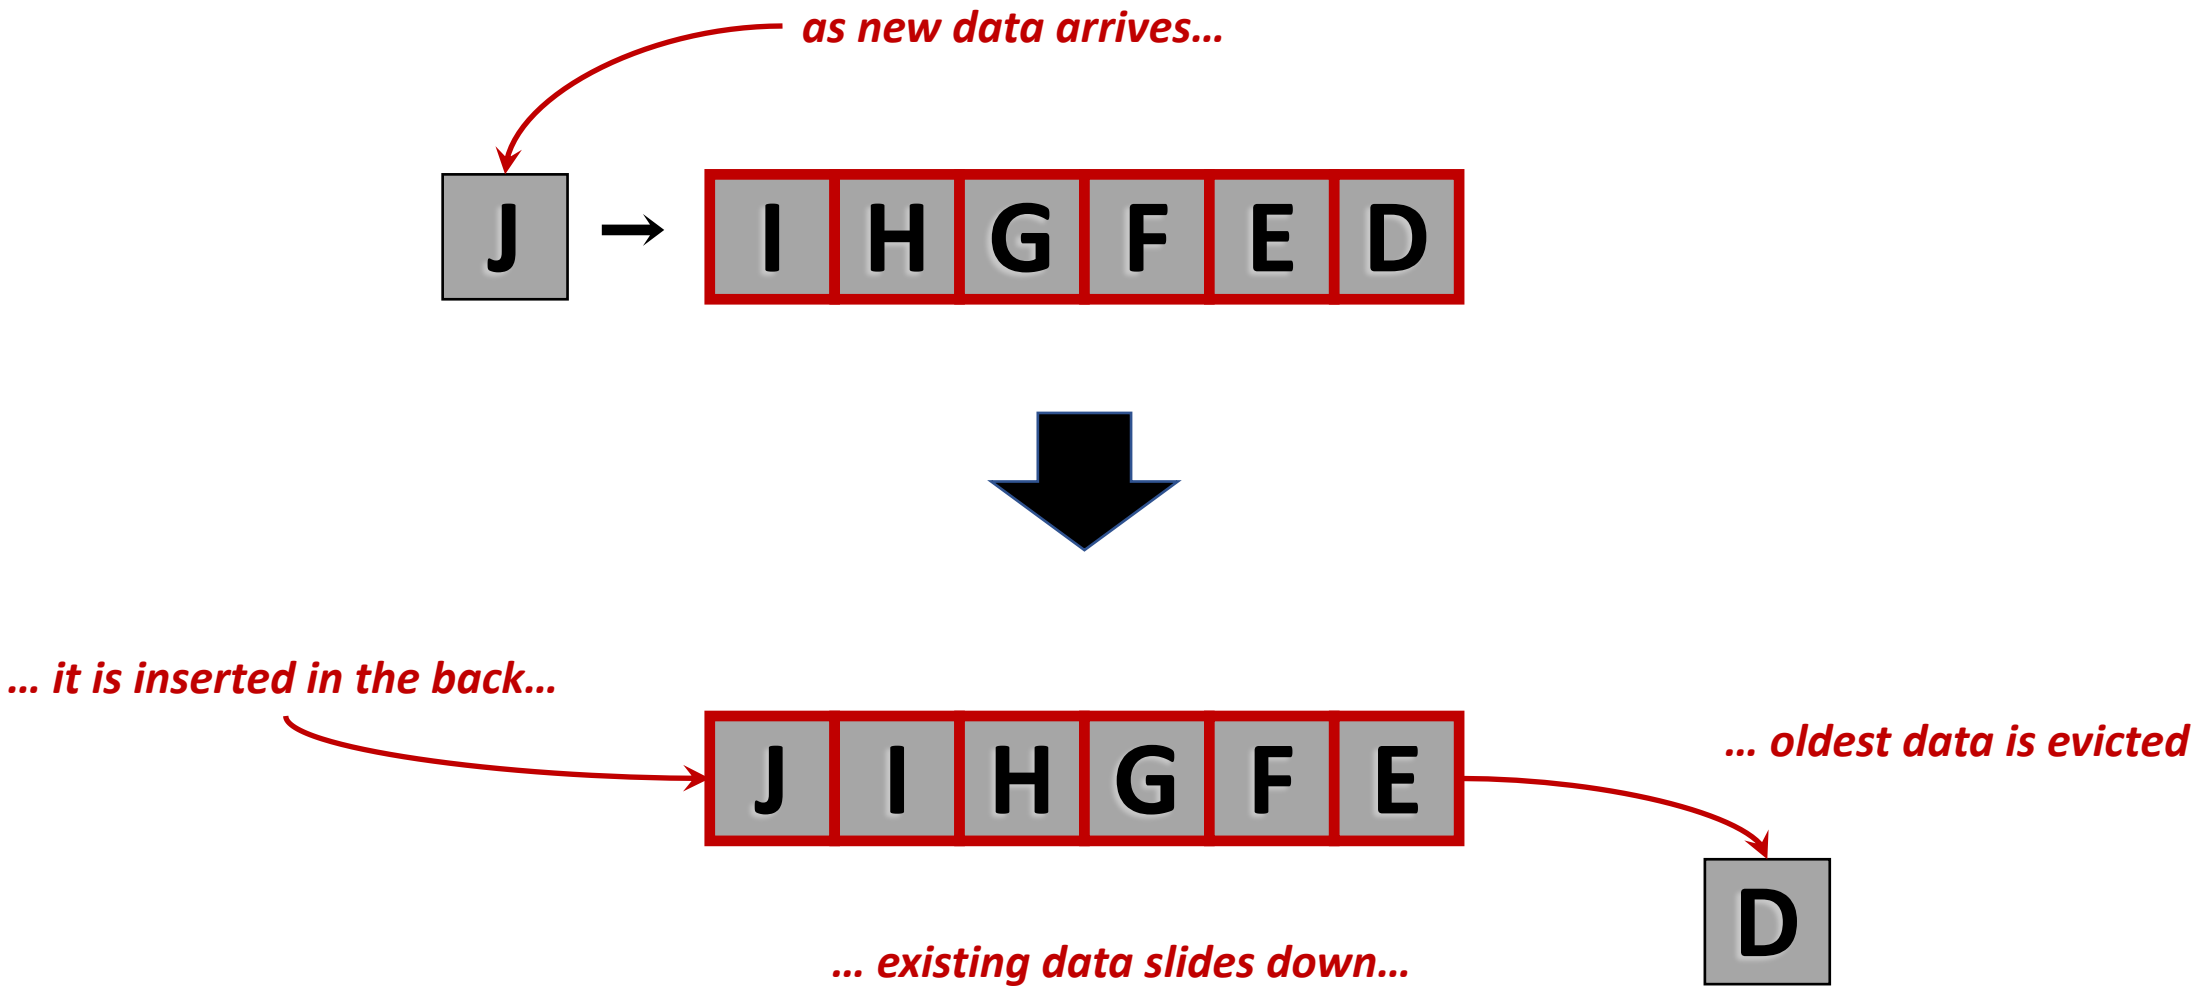

Real data tends to have *timestamps*

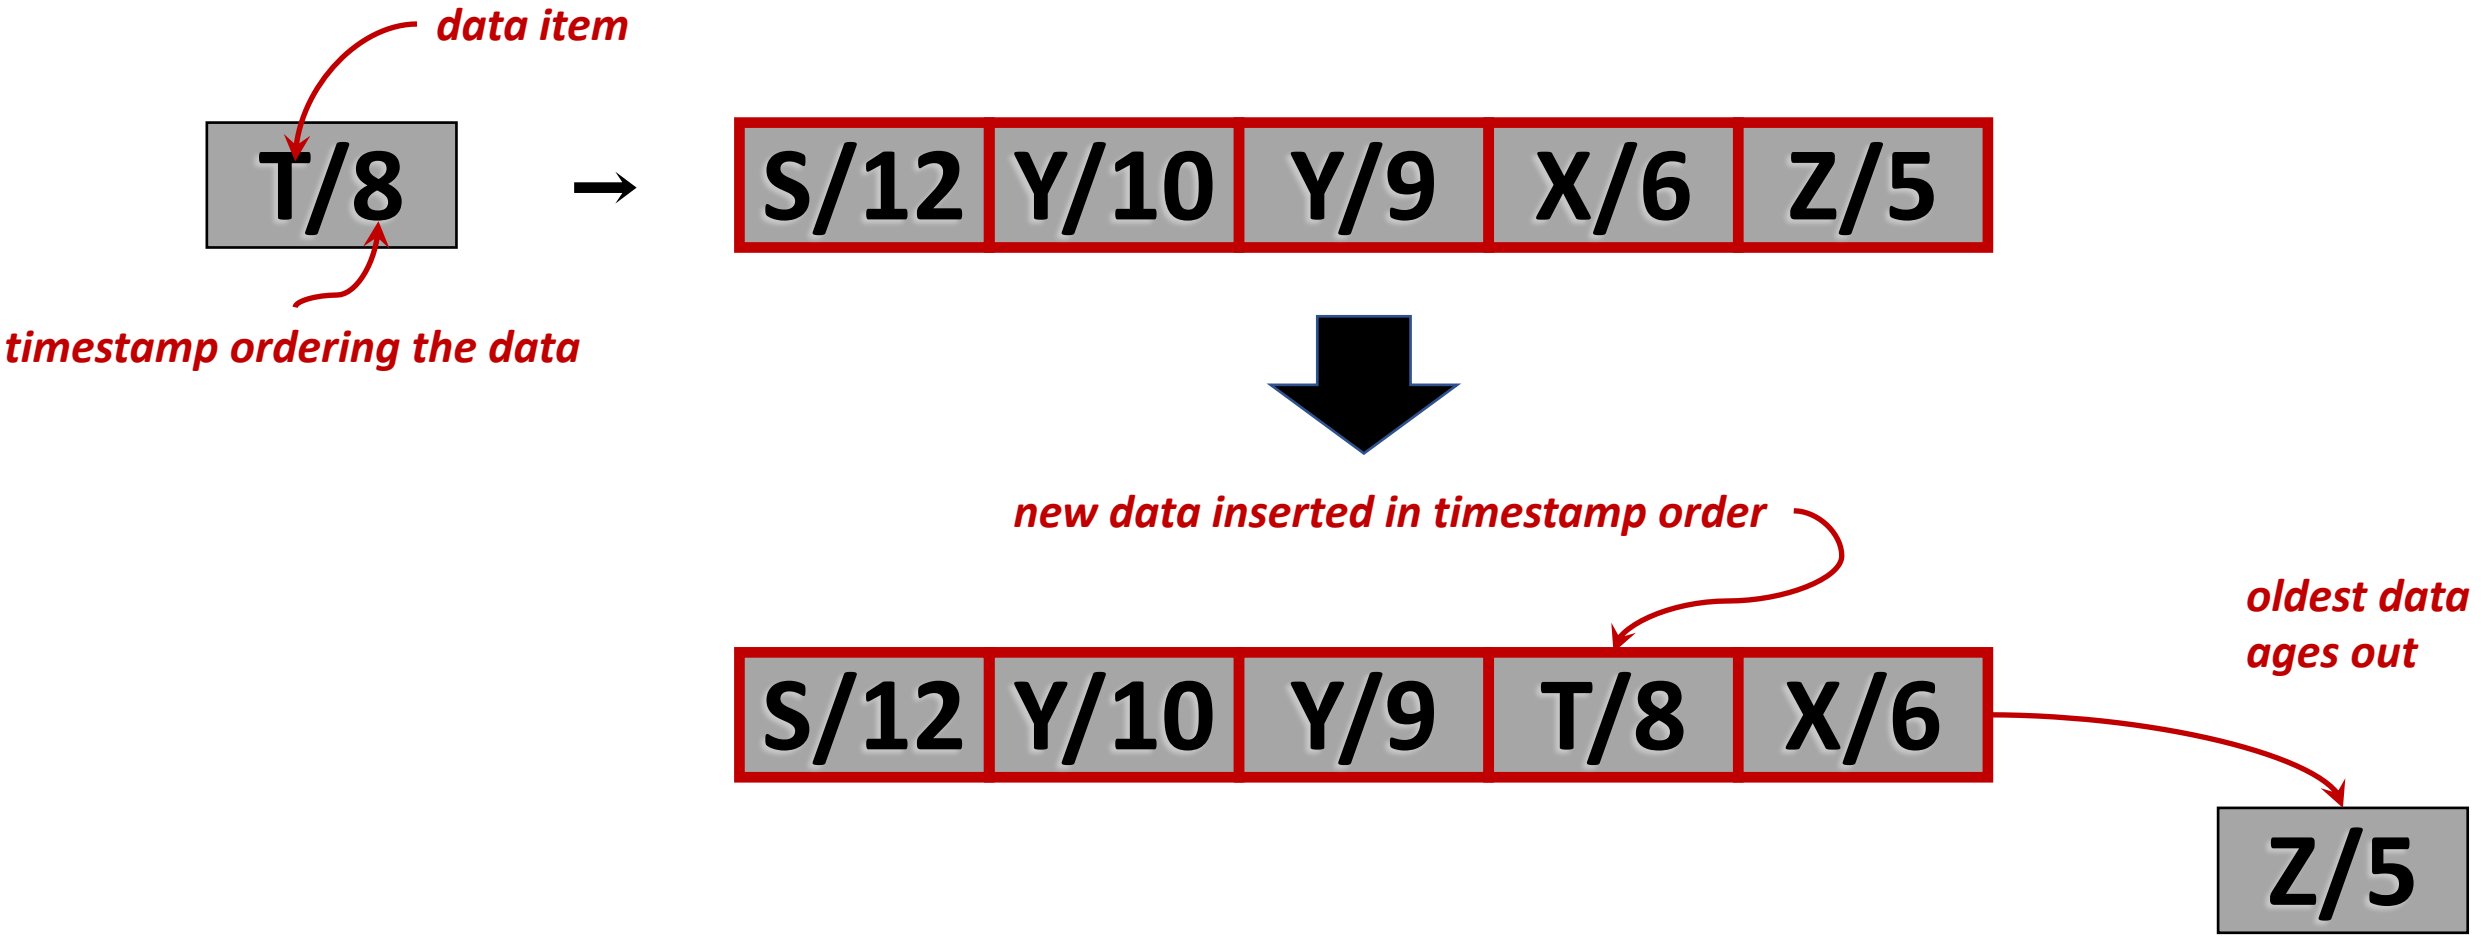

A query on the window is an *aggregation*

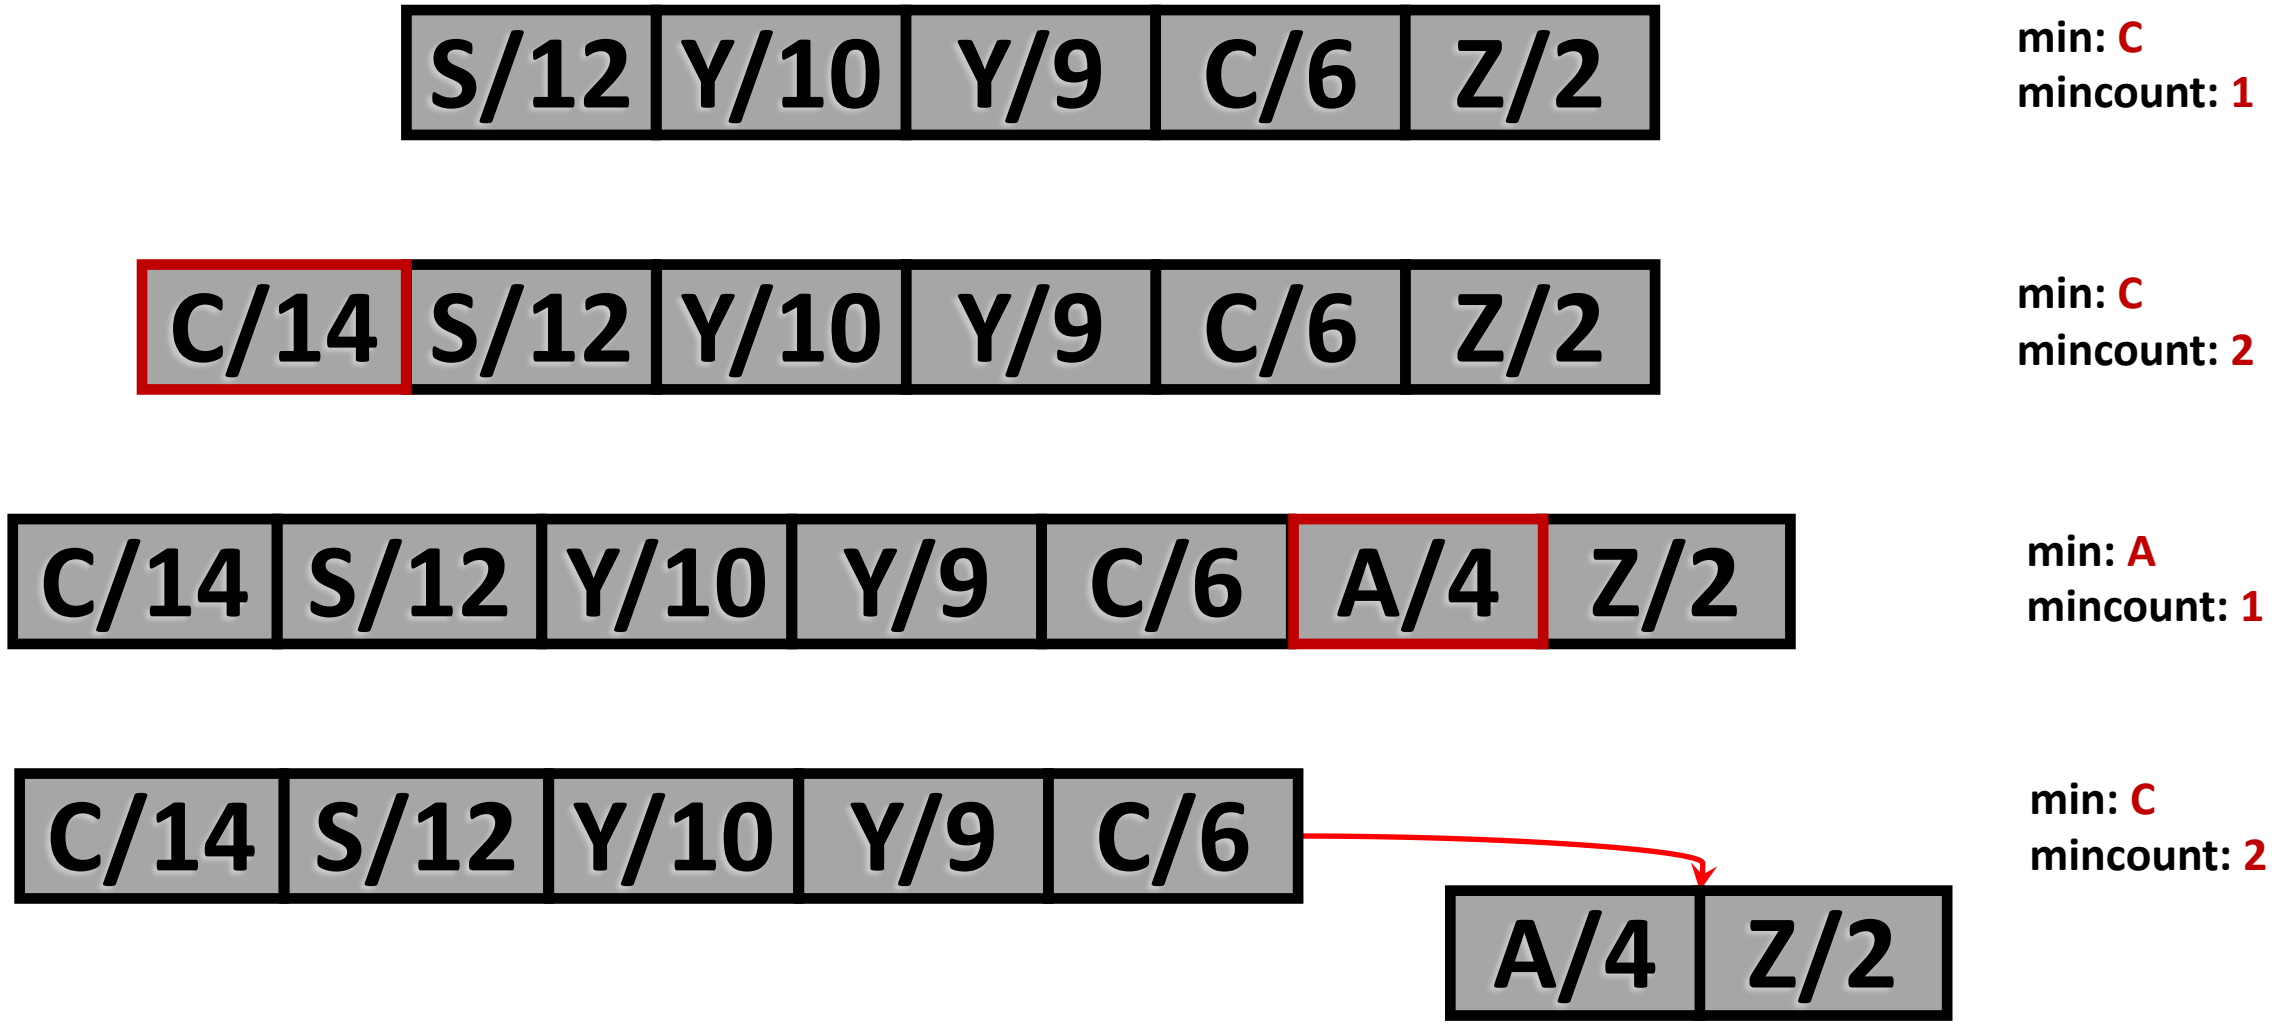

Problem statement

- We want a data structure for a sliding window that can:
  - **insert** data items with timestamps that are out-of-order by distance *d* in amortized *O(*log *d)*, reducing to *O(*1*)* when *d*=0
  - **evict** data items based on timestamps in amortized *O(*log *d)*, reducing to *O(*1*)* when *d*=0
  - **query** the aggregations on the window in worst-case *O(*1*)*
- How?
  - B-Trees! (heavily modified)

FiBA: Finger B-Tree Aggregator

start with a B-Tree modified with aggregates:

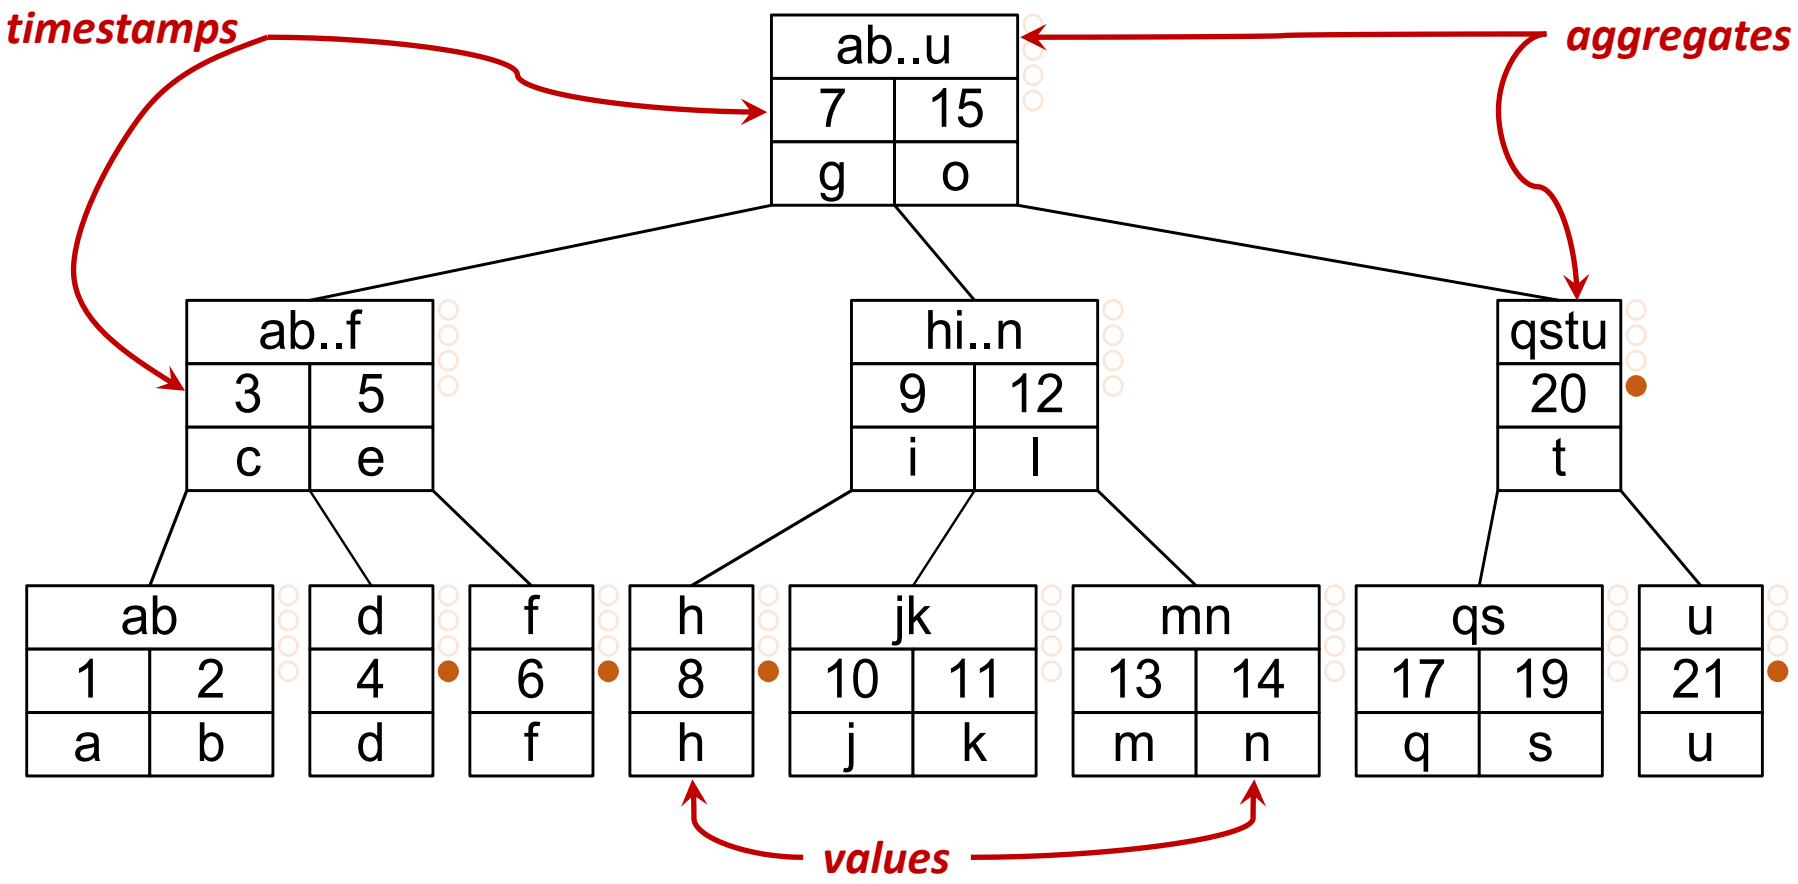

FiBA: Finger B-Tree Aggregator

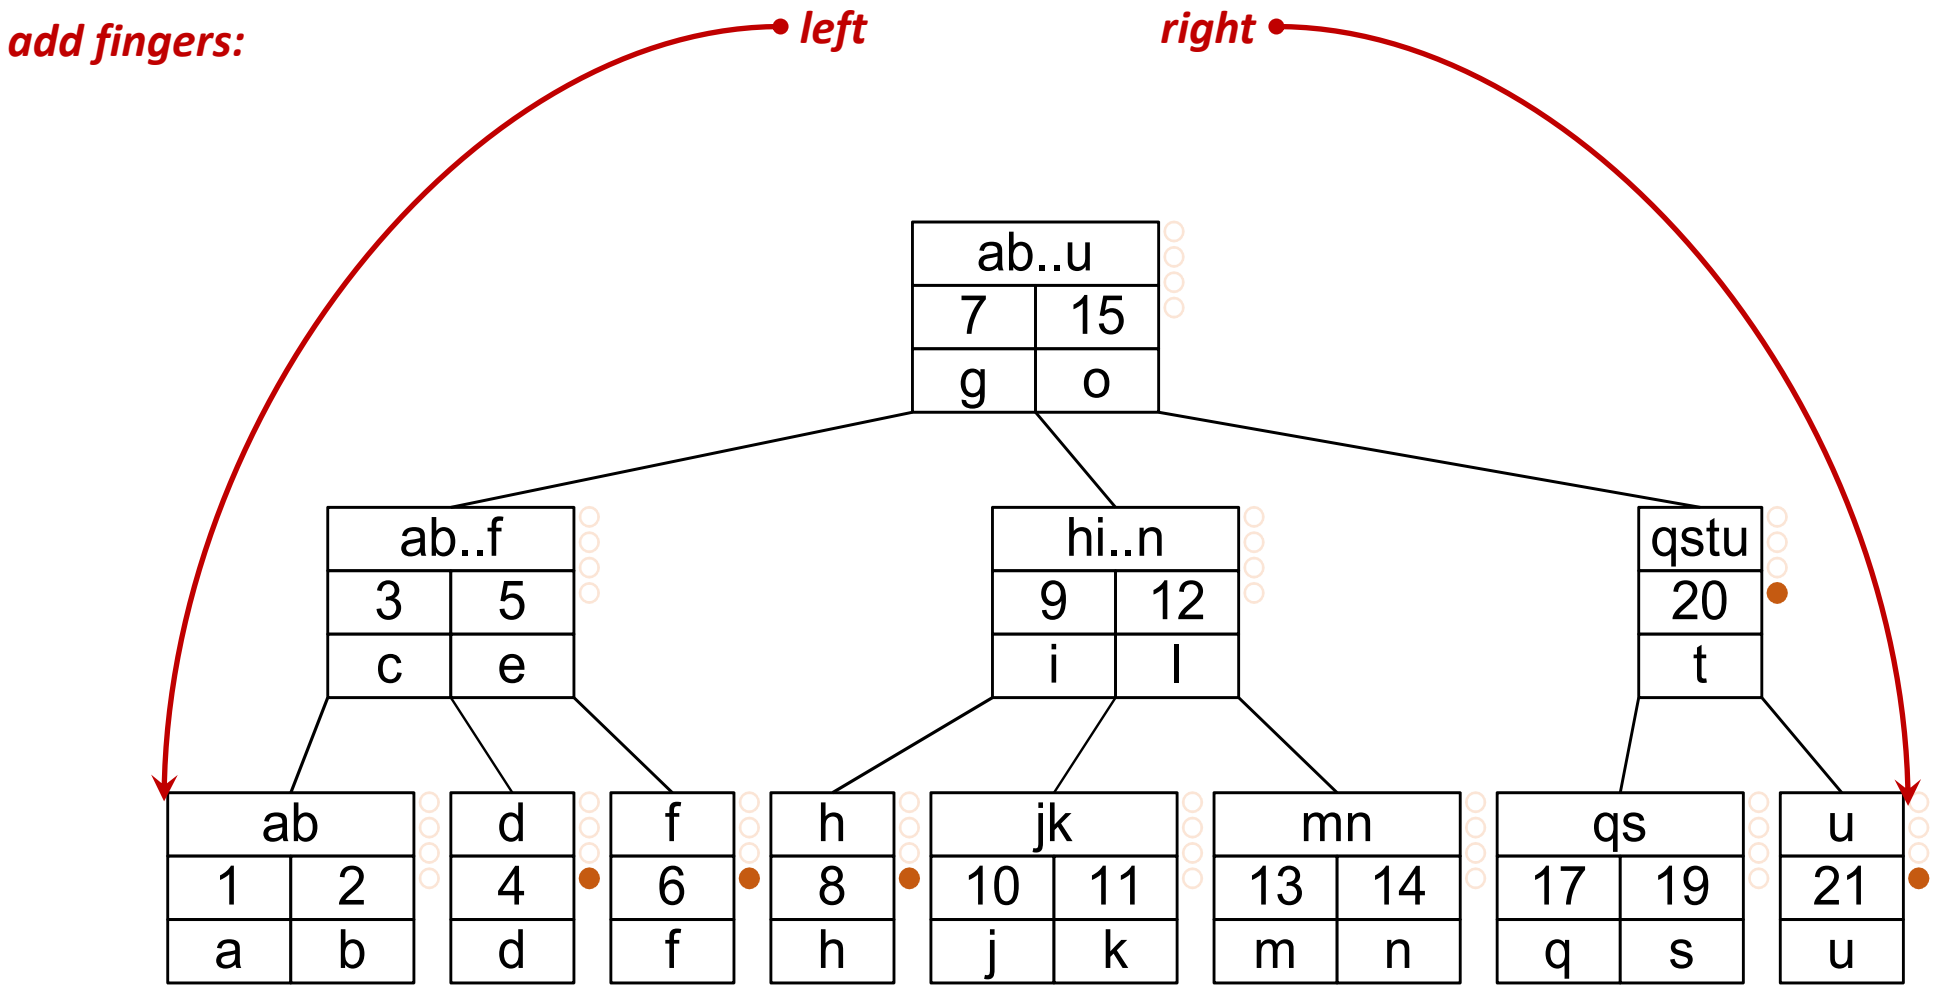

FiBA: Finger B-Tree Aggregator

define position-aware aggregates:

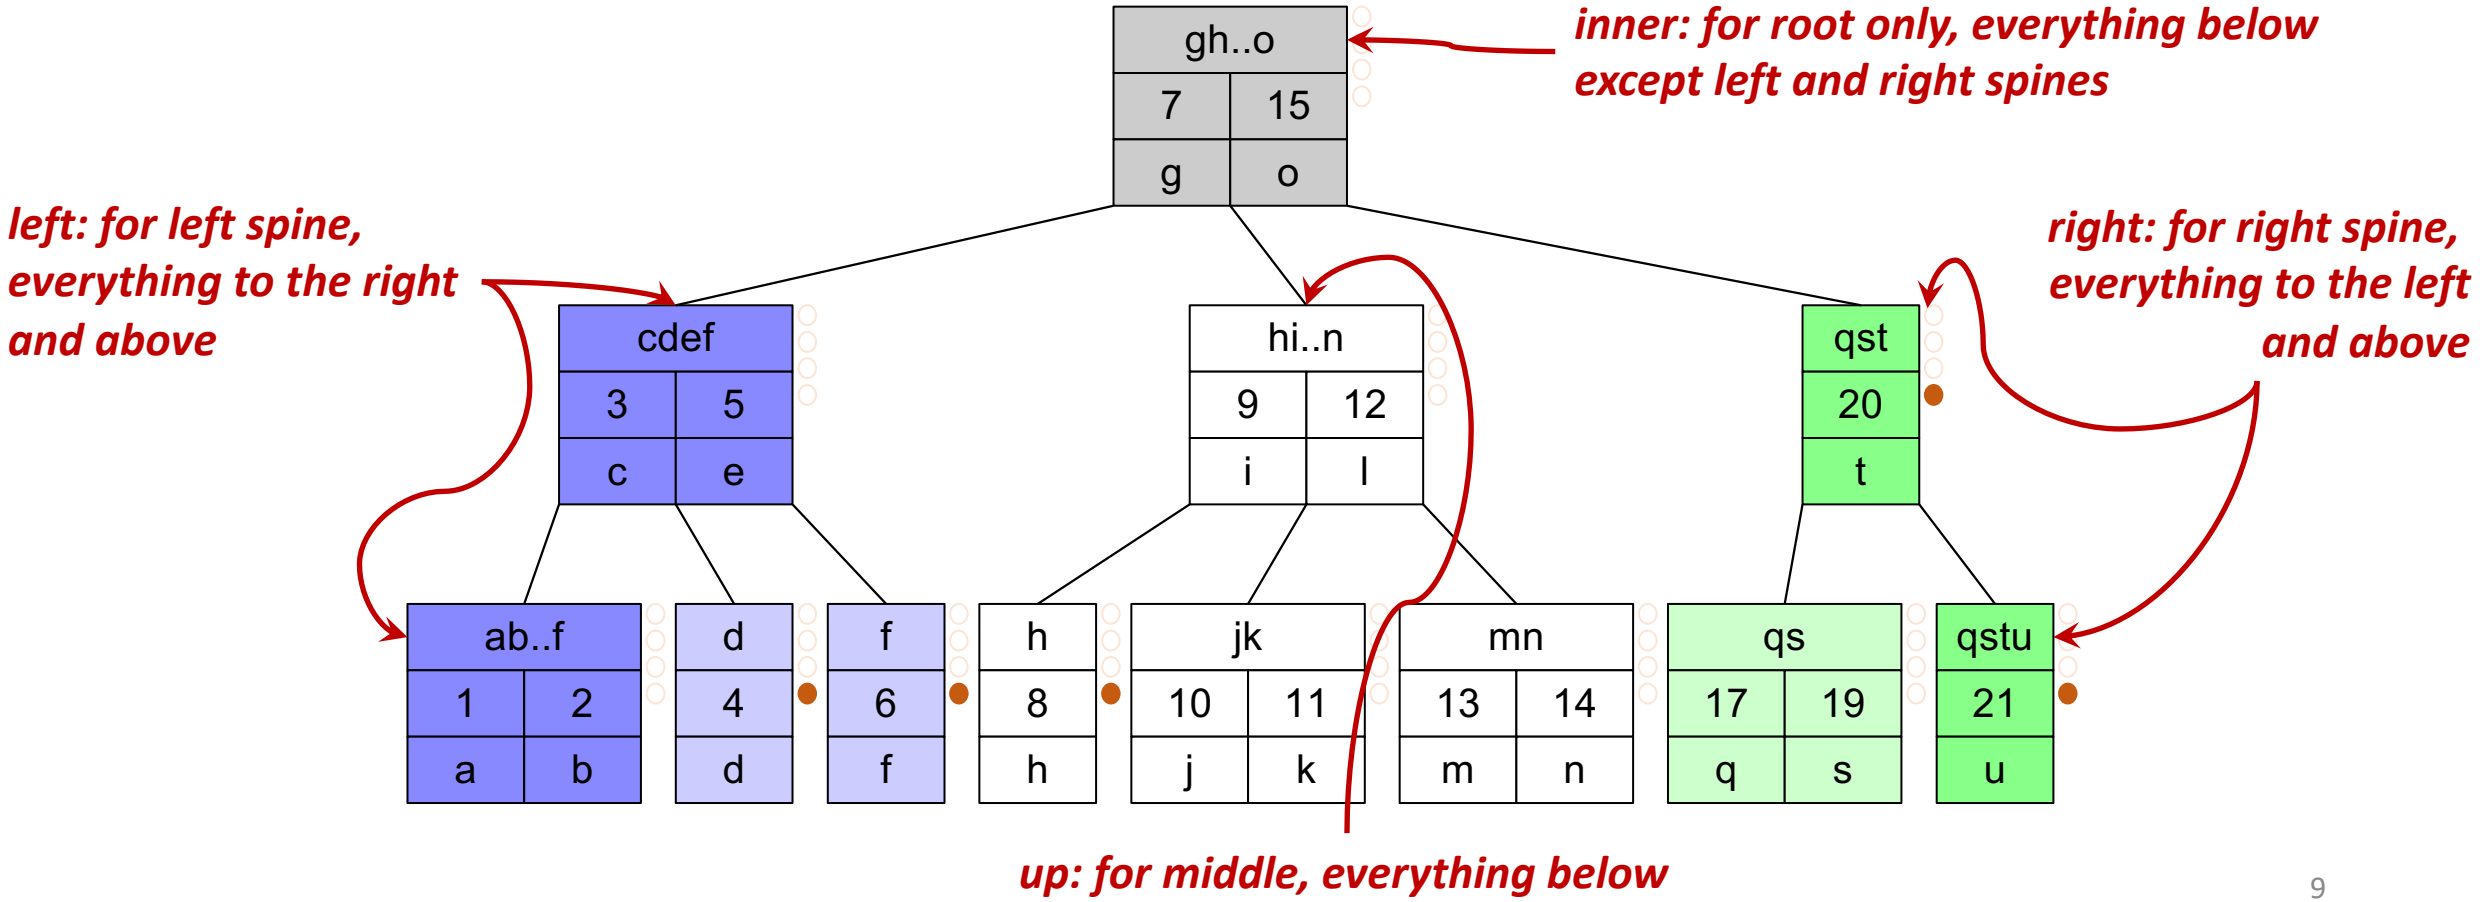

FiBA: Finger B-Tree Aggregator

answer queries by combining left finger, root and right finger:

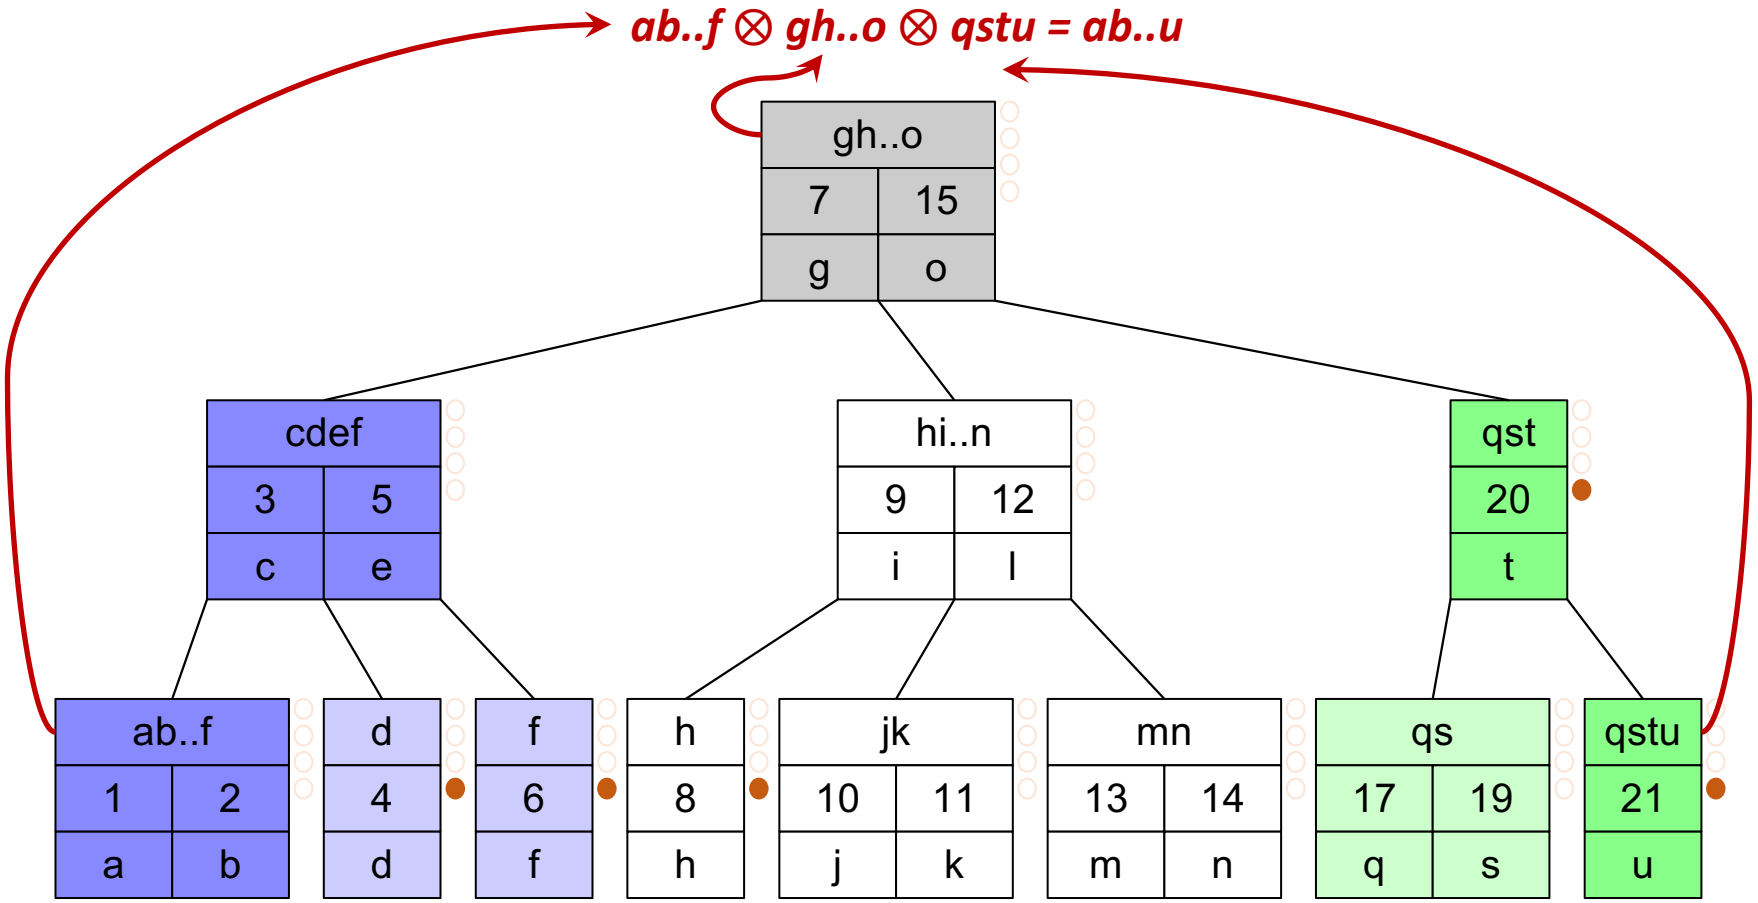

FiBA: Intuition for why it works

- **Fingers** allow *O(1)* time access to oldest and youngest
  - without fingers, searching for a data item would be standard *O(*log *n)*
  - but we are dealing with a time-based window, where we are biased towards inserting at the young end and evicting from the old end
- Specially defined aggregates **shield** sections of the tree
  - updating a value in one section of the tree is unlikely to cause repairs to aggregates elsewhere
- Choice of min and max arity plus lazy splitting and merging **avoids** unnecessary tree rebalancing

Out-of-Order Throughput

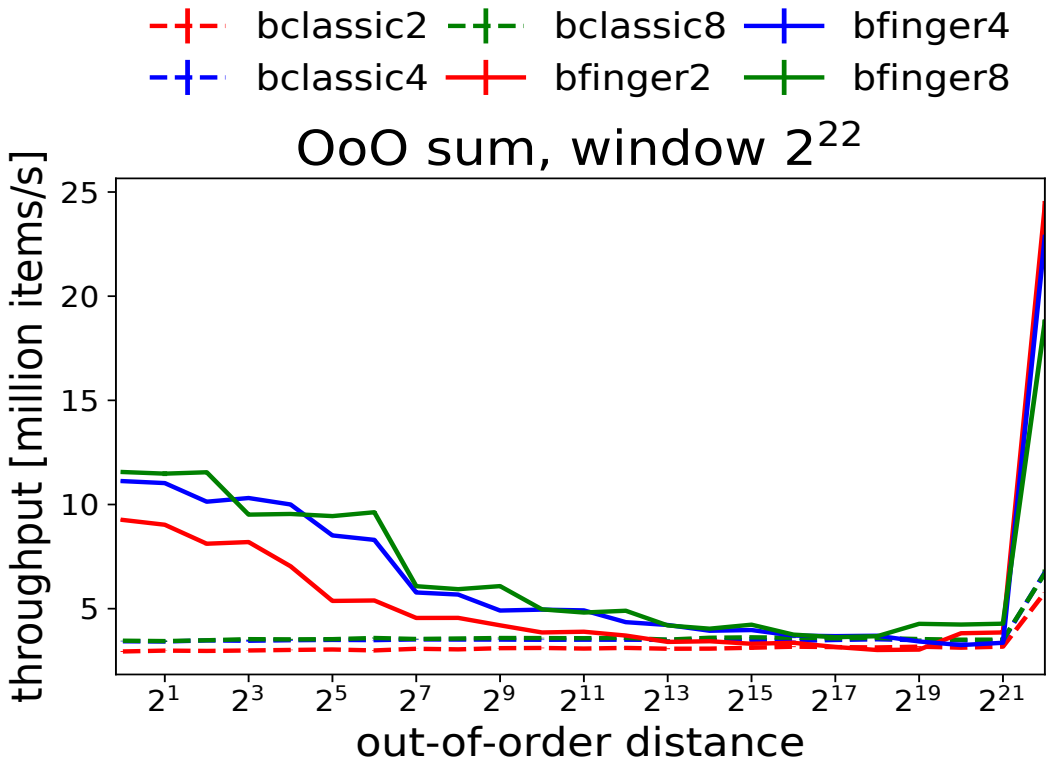

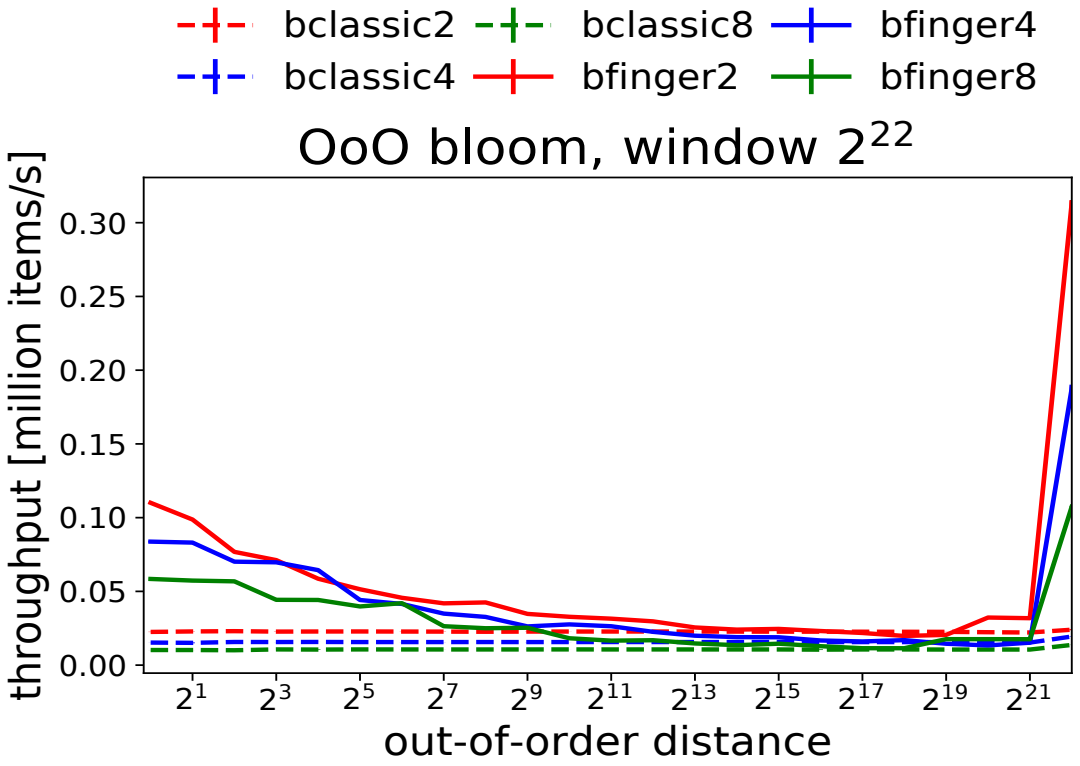

In-Order Latency

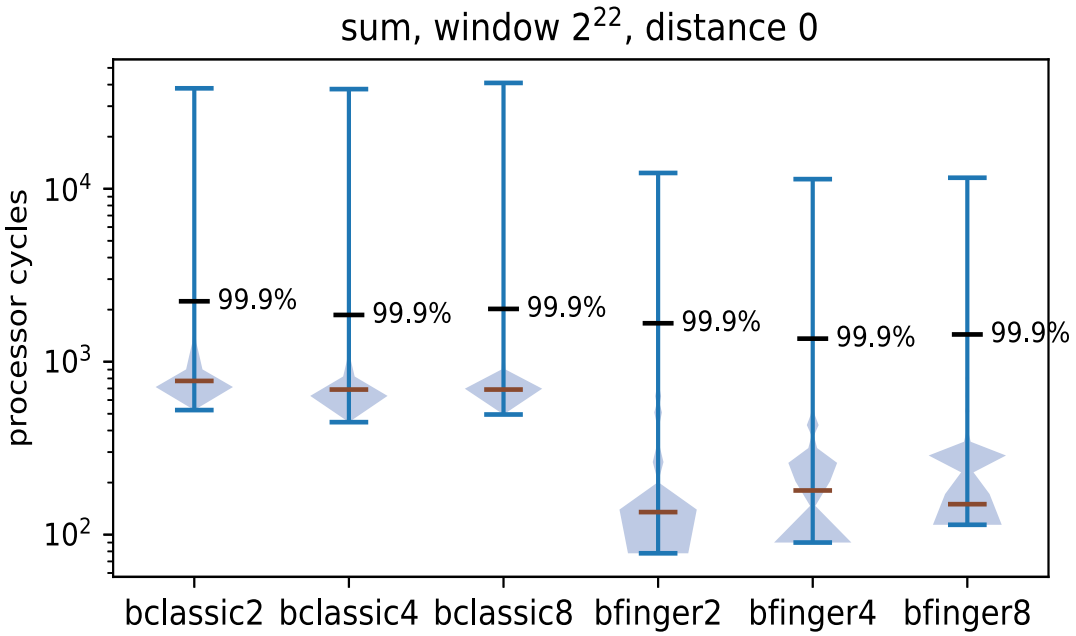

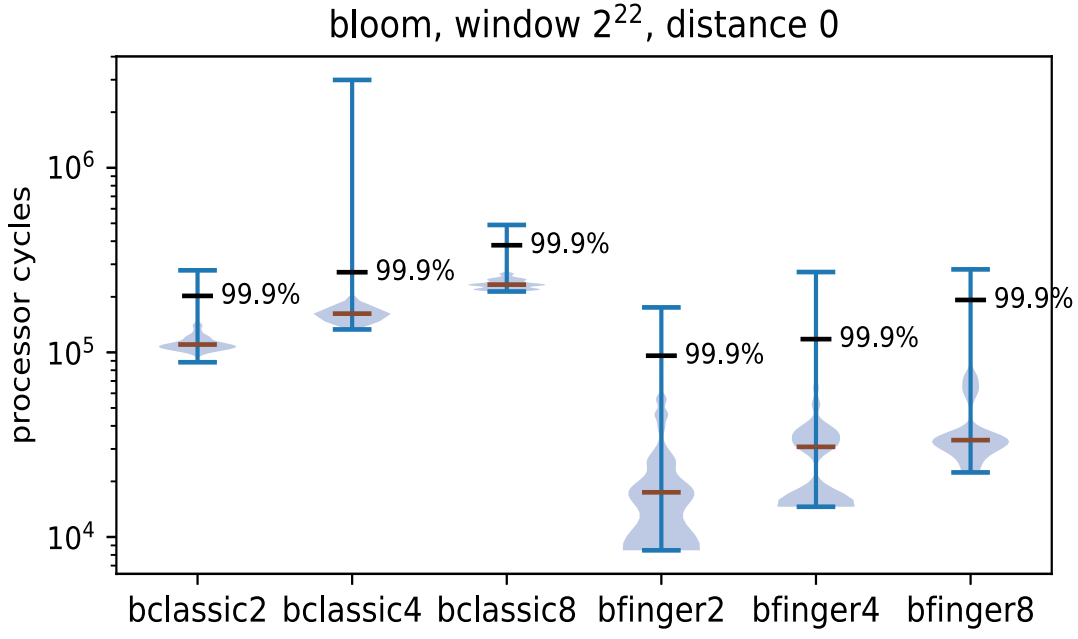

In-Order Throughput

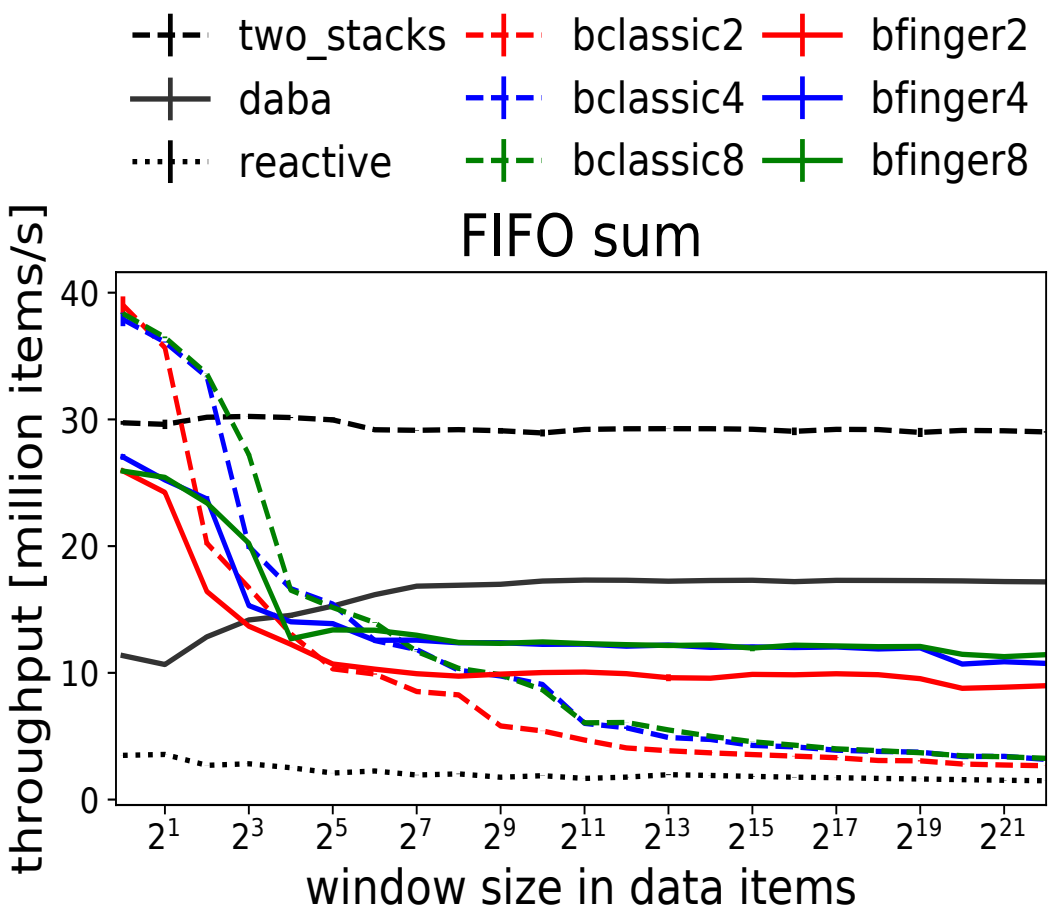

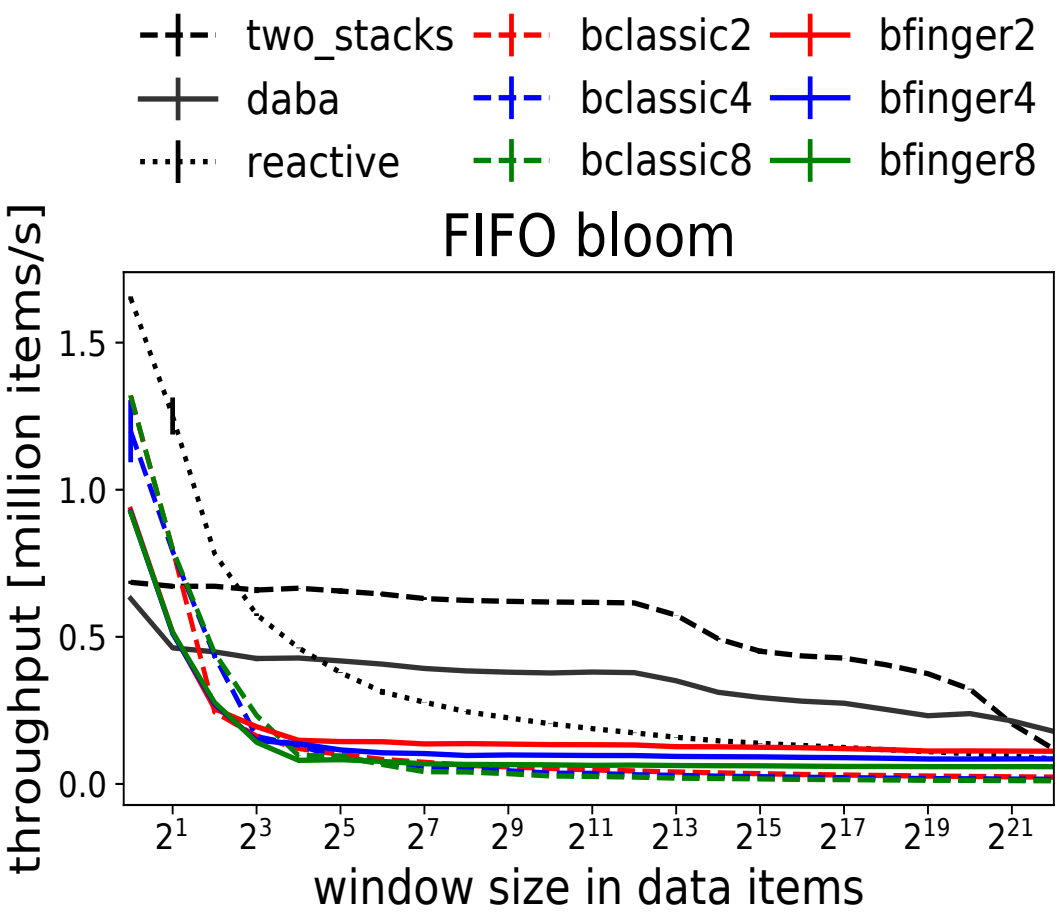

Questions?

Backup

Out-of-Order Latency