# Overlap Interval Partition Join 

University of Zurich ${ }^{\text {UZH }}$<br>Anton Dignös, Michael H. Böhlen, and Johann Gamper<br>\{adignoes, boehlen\}@ifi.uzh.ch, gamper@inf.unibz.it

## GOAL AND APPROACH

- Efficient interval partitioning for the overlap join in valid-time databases
- Find all pairs of tuples with overlapping intervals
- Partition intervals according to position and duration
- Self-adjusting: automatically determine the optimal number of partitions


## Overlap Interval Partitioning - OIP

- Time range is divided into $k$ granules of equal duration
- Partitions are any sequence of contiguous granules


Low $k \Rightarrow$ fewer partition accesses
(less overlapping boxes)


High $k \Rightarrow$ more precise partitions (better fitting boxes)

## Challenges

- Determine partition parameter as a tradeoff between number of false hits and number of partition accesses
- Efficient access structure that allows to omit empty partitions


## Cost Dimensions

False hits: Overhead for tuples that are fetched for $Q$, but are not part of the result.

- CPU cost: identifying and discarding
- IO cost: more data is fetched

Partition accesses: Overhead for fetching and accessing partitions for $Q$.

- CPU cost: search in the access structure
- IO cost: more partially filled blocks

False hits and partition accesses are inversely related.

Constant clustering guarantee: Duration of tuple and partition differs by less than two granules.

## OIP Join

Algorithm

1. Determine parameter $k$ for $\mathcal{O I P}$
2. Partition both input relations using $k$
3. Join tuples within overlapping partitions


## Determining Parameter $k$ for The OIPJoin

Approach: Find $k$ by minimize the overhead cost function $\operatorname{cost}(k)$ w.r.t. $k$ :

cost for partition accesses cost for false hits

## EMPIRICAL EXPERIMENTS



Cost function compared to Runtime



CPU cost 10 cost



CPU cost / IO cost
$k$ adapting to CPU and IO cost


Impact of long tuples


Personnel and file change data


## SUMMARY

Summary

- OIP partitions intervals according to position and duration.
- Long tuples in the datasets do not deteriorate the performance of $\mathcal{O I P}$.
- OIP Join is self-adjusting: $k$ is determined by minimizing the total cost for false hits and partition accesses.
- OIP Join is robust for long tuples.


## Future Work

- Advanced statistics to calculate the number of empty partitions and the reduced average number of partition accesses APA, for instance using histograms.
- Study the maintenance of $\mathcal{O I P}$.
- Refinement of cost function for different buffer replacement strategies.

