

OVERLAP INTERVAL PARTITION JOIN

GOAL AND APPROACH

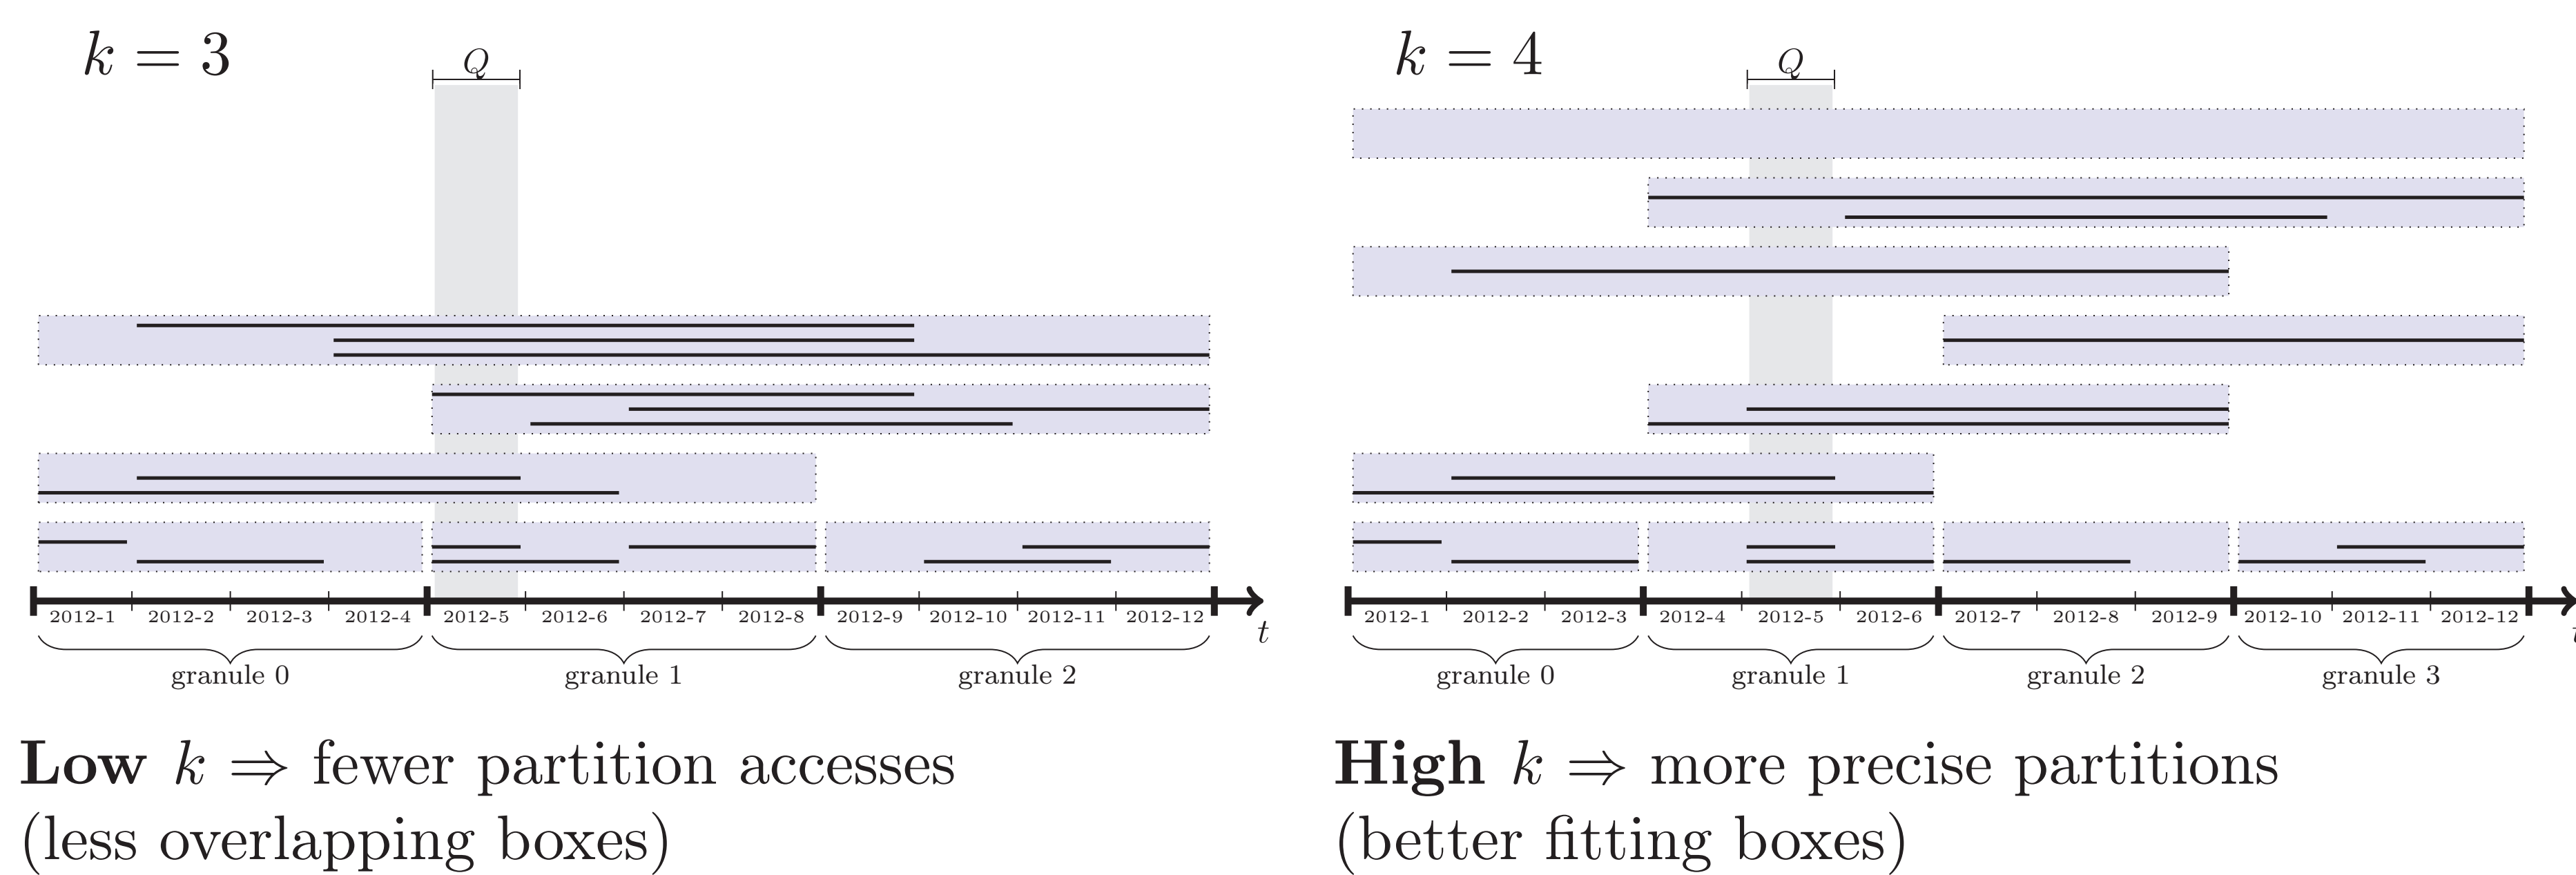
- Efficient **interval partitioning** for the overlap join in valid-time databases
 - Find all pairs of tuples with overlapping intervals
- Partition intervals according to **position** and **duration**
- Self-adjusting: automatically determine the **optimal number of partitions**

CHALLENGES

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

OVERLAP INTERVAL PARTITIONING - OIP

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



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

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.

OIPJOIN

Algorithm

1. Determine parameter k for OIP
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 $cost(k)$ w.r.t. k :

$$cost(k) = |p_r| \cdot \text{APA} \cdot (c_{io} + 2 \cdot c_{cpu}) + |p_r| \cdot n_s \cdot \text{AFR} \cdot \left(\frac{c_{io}}{b} + 2 \cdot \frac{n_r}{|p_r|} \cdot 2 \cdot c_{cpu} \right)$$

partially filled blocks (1 trailing block per partition) search in access structure (2 comparisons in access list)


cost for partition accesses

cost for false hits


more data is fetched (1 false hit within a block) identifying and discarding (2 comparisons per false hit)

Result: $k = f(\text{CPU cost, IO cost, relation sizes})$



EMPIRICAL EXPERIMENTS



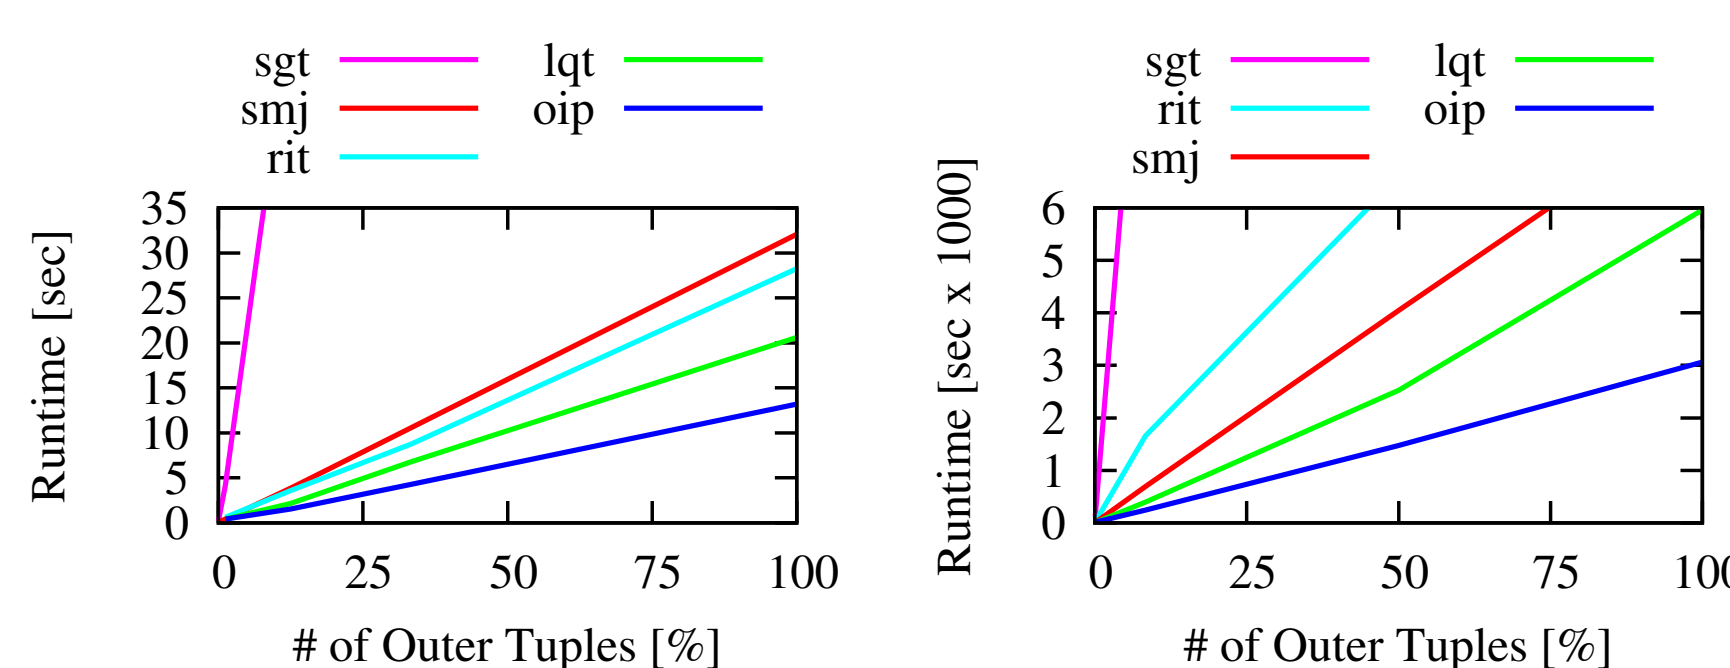
Cost function compared to Runtime



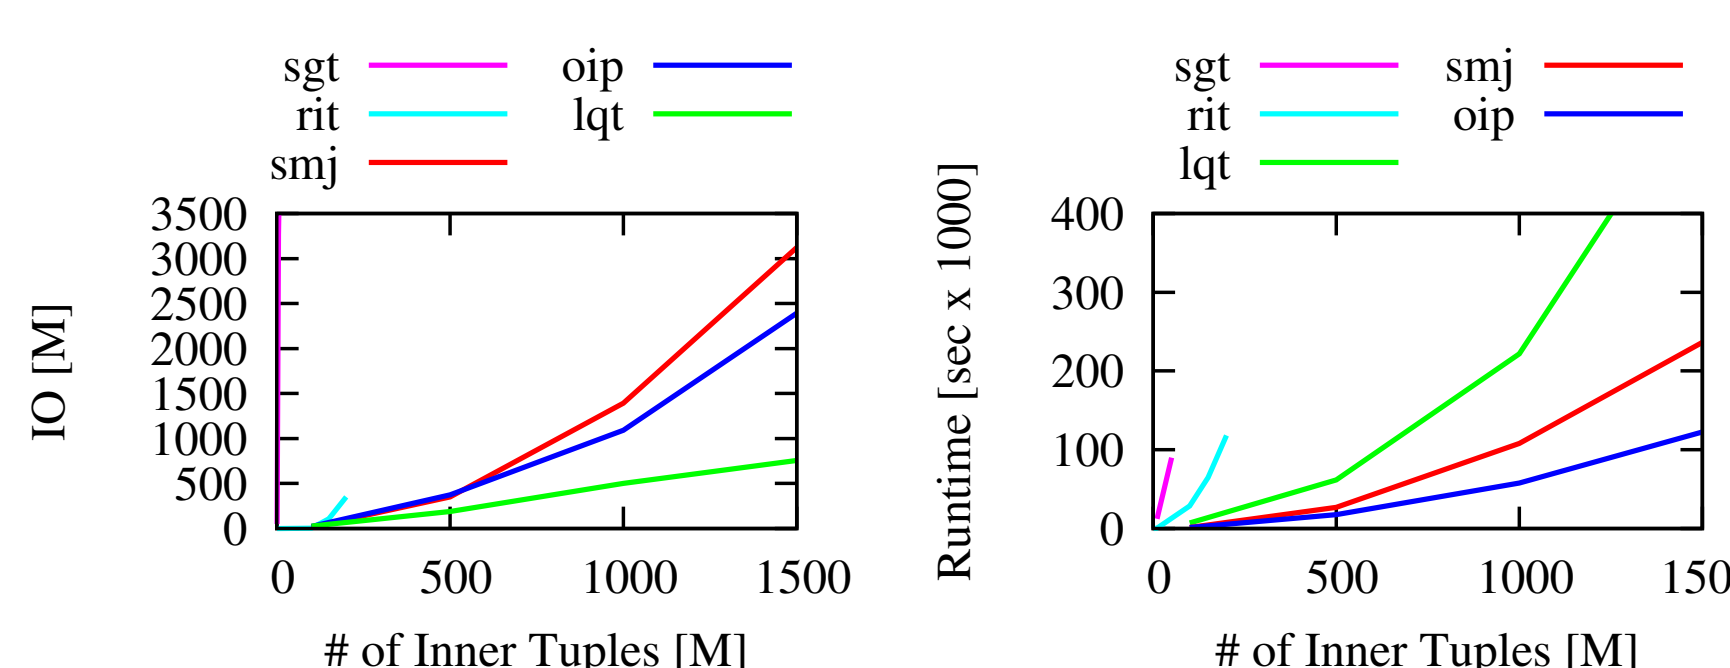
k adapting to CPU and IO cost



Impact of long tuples



Personnel and file change data



Disk Resident Data

SUMMARY

Summary

- OIP partitions intervals according to position and duration.
- Long tuples in the datasets do not deteriorate the performance of OIP .
- OIPJOIN is self-adjusting: k is determined by minimizing the total cost for false hits and partition accesses.
- OIPJOIN 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 OIP .
- Refinement of cost function for different buffer replacement strategies.