

Stab-Forests: Dynamic Data Structures for Efficient Temporal Query Processing

Jelle Hellings

Yuqing Wu



Temporal event-based data

Consider *events* of the form $\langle \text{start-time}, \text{end-time} \rangle$.

Example

- ▶ Process runtime in computer systems.
- ▶ Airline On-Time Performance Data (AOTPD).
- ▶ Civil Unrest Event Data (CUED).

Definition

Let R and S be sets of events. The temporal join $R \bowtie S$ is defined by

$$R \bowtie S = \{ (E_1, E_2) \in R \times S \mid E_1 \cap E_2 \neq \emptyset \}.$$


High-performance temporal joins $R \bowtie S$ via FWDSCAN

- ▶ All data in internal memory
 $\implies R$ and S : *arrays* of events, ordered on time.
- ▶ Cache-friendly traversal of data
 \implies *sequential scan* of R and S .

Proposition


FWDSCAN(R, S) computes $R \bowtie S$ in worst-case $O(|R| + |S| + |output|)$.

Temporal joins via FWDSCAN (example)




$[(R_0, S_0), (R_0, S_1), (R_0, S_2), (R_1, S_0), (R_1, S_1), (R_3, S_2), (R_3, S_3), (R_4, S_3)]$

Temporal joins via FWDSCAN (example)




$[(R_0, S_0), (R_0, S_1), (R_0, S_2), (R_1, S_0), (R_1, S_1), (R_3, S_2), (R_3, S_3), (R_4, S_3)]$

Temporal joins via FWDSCAN (example)




$[(R_0, S_0), (R_0, S_1), (R_0, S_2), (R_1, S_0), (R_1, S_1), (R_3, S_2), (R_3, S_3), (R_4, S_3)]$

Temporal joins via FWDSCAN (example)




$[(R_0, S_0), (R_0, S_1), (R_0, S_2), (R_1, S_0), (R_1, S_1), (R_3, S_2), (R_3, S_3), (R_4, S_3)]$

Temporal joins via FWDSCAN (example)




$[(R_0, S_0), (R_0, S_1), (R_0, S_2), (R_1, S_0), (R_1, S_1), (R_3, S_2), (R_3, S_3), (R_4, S_3)]$

Temporal joins via FWDSCAN (example)




$[(R_0, S_0), (R_0, S_1), (R_0, S_2), (R_1, S_0), (R_1, S_1), (R_3, S_2), (R_3, S_3), (R_4, S_3)]$

Temporal joins via FWDSCAN (example)




$[(R_0, S_0), (R_0, S_1), (R_0, S_2), (R_1, S_0), (R_1, S_1), (R_3, S_2), (R_3, S_3), (R_4, S_3)]$

Temporal joins via FWDSCAN (example)




$[(R_0, S_0), (R_0, S_1), (R_0, S_2), (R_1, S_0), (R_1, S_1), (R_3, S_2), (R_3, S_3), (R_4, S_3)]$

Temporal joins via FWDSCAN (example)



$[(R_0, S_0), (R_0, S_1), (R_0, S_2), (R_1, S_0), (R_1, S_1), (R_3, S_2), (R_3, S_3), (R_4, S_3)]$

Temporal joins via FWDSCAN (example)



$[(R_0, S_0), (R_0, S_1), (R_0, S_2), (R_1, S_0), (R_1, S_1), (R_3, S_2), (R_3, S_3), (R_4, S_3)]$

Question

Can we prevent inspection of R_2 ?

The SKIPJOIN algorithm

Algorithm SKIPJOIN(R, S):

```
1:  $i, j := 0, 0$ 
2: while  $i < |R|$  and  $j < |S|$  do
3:   if  $R[i].start \leq S[j].start$  then
4:     if  $S[j].start \leq R[i].end$  then       $\triangleright$  Normal Join (e.g., Forward Scan).
5:       Join  $R[i]$  with  $S[j\dots]$ 
6:        $i := i + 1$ 
7:     else                                   $\triangleright$  Skip Events.
8:        $(i, L) := \text{STAB}(R[i\dots], S[j].start)$ 
9:       For each event  $E \in L$ , join  $E$  with  $S[j\dots]$ 
10:    else analogous (swap roles of  $R$  and  $S$ )
```

$$\text{STAB}(R, t) = \{E \in R \mid E.start \leq t \leq E.end\}$$

High-performance sequences of stab queries

Problem


Cannot efficiently query an array A for all events active at t .

Solution


Stab-forest: index-structure supporting sequences of stab-queries.

- ▶ A fully-balanced binary search tree to minimize size.
- ▶ Event-based augmentations via compact static arrays.
- ▶ Efficiently supports append-only operations.


Stab-forests: Collections of stab-trees



Stab-forests: Collections of stab-trees




Stab-forests: Querying a stab-tree




Query: “MULTISTAB(\mathcal{S} , $[0, 2, 7]$)”.

Stab-forests: Querying a stab-tree




Query: “MULTISTAB(\mathcal{S} , $[0, 2, 7]$)”.

Stab-forests: Querying a stab-tree




Stab-forests: Querying a stab-tree




Query: “MULTISTAB(\mathcal{S} , $[0, 2, 7]$)”.

Stab-forests: Cache-friendly querying a stab-tree




Query: “MULTISTAB(\mathcal{S} , [0, 2, 7])”.

Stab-forests: Cache-friendly querying a stab-tree




Query: “MULTISTAB(\mathcal{S} , $[0, 2, 7]$)”.

Stab-forests: Cache-friendly querying a stab-tree



Query: “MULTISTAB(\mathcal{S} , $[0, 2, 7]$)”.

Stab-forests: Cache-friendly querying a stab-tree



Query: “MULTISTAB(\mathcal{S} , $[0, 2, 7]$)”.

SKIPJOIN with stab-forests: Results

Theorem

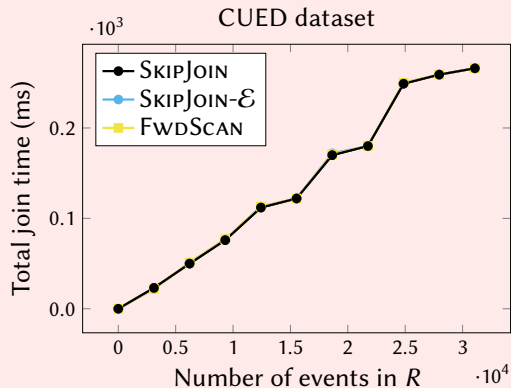
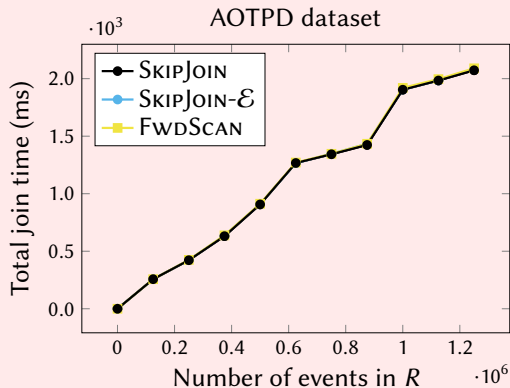
A stab-forest indexing L

- ▶ *can be stored in worst-case $O(|L|)$ space.*
- ▶ *can be constructed in $O(|L| \log|L|)$.*
- ▶ *can be appended-to in amortized $O(\log|L|)$.*
- ▶ *can be stabbed for timestamps ϕ in $O(\min(|\phi| \log|L|, |\phi| + |L|) + |\text{output}|)$.*

Theorem

SKIPJOIN(R, S) computes $R \bowtie S$ in worst-case $O(M(R, S) + M(S, R) + |\text{output}|)$, with $M(A, B)$ the cost of stabbing A with $|B|$ timestamps.

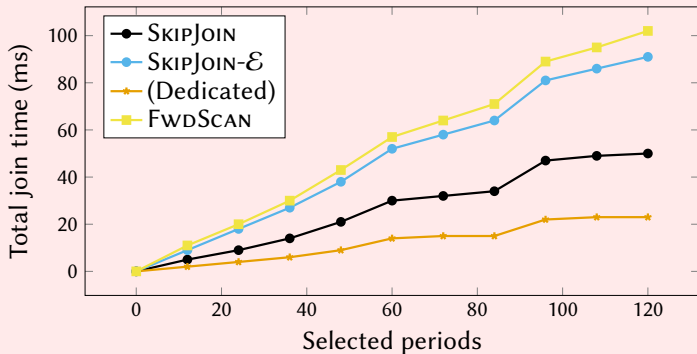
Self-joining real-world datasets



- ▶ 61,100,539 events in 10 years.
- ▶ Shortest: 0 min. Longest: 1350 min.
- ▶ Used: 2,500,000 events.

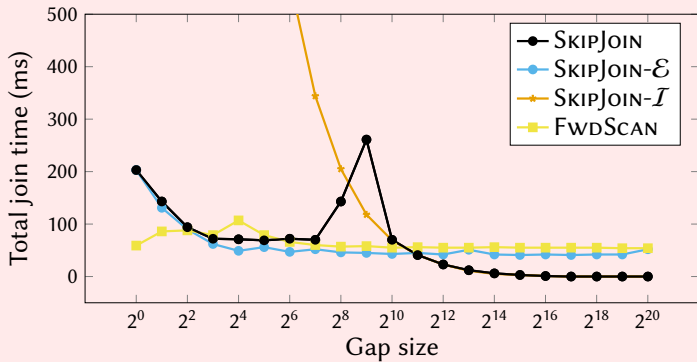
- ▶ 62,141 events in 60 years.
- ▶ Shortest: 0 d. Longest: 18 407 d.

Sparse joins and dedicated multi-window-queries




Query: *“select all events on the 7-th day from the first n months in AOTPD”.*

The behavior of SKIPJOIN

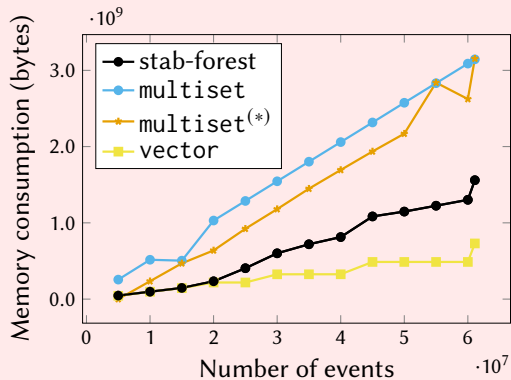
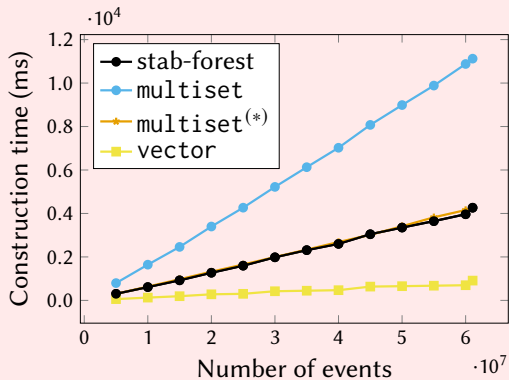


Query: “ $R \bowtie S$ with



(gap size 3)”.

The costs of stab-forests



Conclusion

SKIPJOIN: (windowed) temporal joins on skewed datasets.

Stab-forests: append-only index for sequences of stab-queries.

Future Work

- ▶ Using stab-forests for other temporal operations.
- ▶ Block-based external-memory designs based on the append-only stab-forests.
- ▶ Parallel and distributed high-performance temporal joins.

<https://jhellings.nl/>