

# TreeTracker Join: Simple, Optimal, Fast

ANONYMOUS AUTHOR(S)

We present a novel linear-time acyclic join algorithm, TreeTracker Join (TTJ). The algorithm can be understood as the pipelined binary hash join with a simple twist: upon a hash lookup failure, TTJ resets execution to the binding of the tuple causing the failure, and removes the offending tuple from its relation. Compared to the best known linear-time acyclic join algorithm, Yannakakis's algorithm, TTJ shares the same asymptotic complexity while imposing lower overhead. Further, we prove that when measuring query performance by counting the number of hash probes, TTJ will match or outperform binary hash join on the same plan. This property holds independently of the plan and independently of acyclicity. We are able to extend our theoretical results to cyclic queries by introducing a new hypergraph decomposition method called tree convolution. Tree convolution iteratively identifies and contracts acyclic subgraphs of the query hypergraph. The method avoids redundant calculations associated with tree decomposition and may be of independent interest. Empirical results on TPC-H, the Join Order Benchmark, and the Star Schema Benchmark demonstrate favorable results.

## ACM Reference Format:

Anonymous Author(s). 2025. TreeTracker Join: Simple, Optimal, Fast. 1, 1 (December 2025), 30 pages. <https://doi.org/10.1145/nnnnnnnnnnnnnn>

## 1 Introduction

Yannakakis [38] was the first to describe a linear-time join algorithm (hereafter YA) running in time  $O(|\text{IN}| + |\text{OUT}|)$ , where  $|\text{IN}|$  is the input size and  $|\text{OUT}|$  is the output size. In principle, this is the best asymptotic complexity one can hope for, because in most cases the algorithm must read the entire input and write the entire output. However, virtually no modern database systems implement YA. A major factor is its high overhead. Prior to executing the join, YA performs two passes over the input relations, using semijoins to reduce the input size. The reduction is lossless and enables optimally joining the reduced relations. Since the cost of a semijoin is proportional to the size of its arguments, this immediately incurs a  $2\times$  overhead in the input size. An improved version of YA [2] achieves the same result in one semijoin pass, but the overhead of this pass remains. Another practical challenge is that YA is “too different” from traditional binary join algorithms, making it difficult to integrate into existing systems. For example, the efficiency of YA critically depends on a query’s *join tree* which is different from the query plan used by binary joins<sup>1</sup>. Where there is a wealth of techniques to optimize query plans for binary joins, little is known about cost-based optimization of join trees for YA.

In this paper, we propose a new linear-time join algorithm called TreeTracker Join (TTJ), inspired by the TreeTracker algorithm [15] in Constraint Satisfaction. TTJ can be understood as the traditional binary hash join with a twist: when a hash lookup fails, backtrack to the tuple causing the failure, and remove that tuple from its relation. The backtracking points depend only on the query, not the data, and are determined by the query compiler prior to query execution. The execution deviates

---

<sup>1</sup>By *join tree* we mean (hyper-)tree decomposition of hypertree width 1, not the tree of binary join operators commonly seen in relational algebra query plans.

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM XXXX-XXXX/2025/12-ART

<https://doi.org/10.1145/nnnnnnnnnnnnnn>

<pre> 1  for i,x in R:           for i,x in R:           for i,x in R:           1 2    for y,j in S[x]:       for y,j in S[x]:       for y,j in S[x]:       2 3      for k in T[y]:       for k in T[y]:       for k in T[y]:       3 4      # lines left blank   if U[y] is None:      if U[y] is None:      4 5      # intentionally      break                  S[x].del((y,j)); break 5 6      for l in U[y]:       for l in U[y]:       for l in U[y]:       6 7        print(x,y,i,j,k,l)  print(x,y,i,j,k,l)  print(x,y,i,j,k,l)  7 </pre>	<p>(a) Binary join</p>	<p>(b) Backjumping</p>	<p>(c) Tuple deletion</p>
--	------------------------	------------------------	---------------------------

Fig. 1. Instantiation of binary hash join on example 1.1, with backjumping, and with tuple deletion.  $S[x]$  denotes a hash lookup into the hash table for  $S$  using  $x$  as key.

from binary hash join only when a dangling tuple is detected and deleted. Thus the tuple is excluded from any computation going forward. Hence, when using identical query plans TTJ is guaranteed to match or outperform binary hash join (Section 4.2).

The following example illustrates the main ideas of TTJ.

*Example 1.1.* Consider the natural join of the relations  $R(i, x)$ ,  $S(x, y, j)$ ,  $T(y, k)$ , and  $U(y, l)$ , where we use  $R(i, x)$  to denote that the schema of  $R$  is  $\{i, x\}$ . The set  $\{1, \dots, N\}$  is denoted by  $[N]$ . Let the relations be defined as follows:

$$R = \{(i, 1) \mid i \in [N]\} \quad S = \{(1, 1, j) \mid j \in [N]\} \quad T = \{(1, k) \mid k \in [N]\} \quad U = \{(0, l) \mid l \in [N]\}$$

Observe  $U$  shares no common  $y$ -values with  $S$  or  $T$ , making the query result empty. We'll first consider execution with binary hash join, the foundation of our algorithm. Assume the optimizer produces a left-deep join plan  $((R \bowtie S) \bowtie T) \bowtie U$ . The execution engine builds hash tables for  $S$ ,  $T$ , and  $U$ , mapping each  $x$  to  $(y, j)$  values in  $S$ ,  $y$  to  $k$  values in  $T$ , and  $y$  to  $l$  values in  $U$ . Figure 1a<sup>2</sup> illustrates the basis of the execution. For each  $(i, x)$  tuple in  $R$ , the hash table for  $S$  is probed for the  $(y, j)$  values.  $T$  is probed with each pair  $(y, j)$  to determine the  $k$  values. This repeats for  $U$ . Although the query produces no output, the execution takes  $\Omega(N^3)$  time because it first computes the join of  $R$ ,  $S$ , and  $T$ . A closer look at the execution reveals the culprit: when the lookup on  $U$  produces no result, (line 6), the algorithm continues to the next iteration of the loop over  $k$  values, (line 3). The same value of  $y$  is used to probe into  $U$  again! To address this, **the first key idea of TTJ is to backjump<sup>3</sup> to the level causing the probe failure.** For clarity we use `break` to represent the backjump, as shown in Figure 1b. When probing  $U$  with the key value  $y$  fails to return a result, we break out of the current loop over  $k$  and continue to the next iteration of the second loop level. That is because the unsuccessful lookup key value of  $y$  is assigned at that level. That next iteration retrieves new  $y, j$  values, skipping over iterations for  $k$  values that are doomed to fail. With this optimization, the execution finishes in  $O(N^2)$  time, as it still needs to compute the join of  $R$  and  $S$ . To improve the performance further: **the second key idea of TTJ is to delete the tuple causing the probe failure.** This is shown in Figure 1c: after the probe failure, the offending tuple  $(x, y, j)$  is removed from the  $S$  hash table. This is safe to do, because that  $y$  value will always fail to join with any tuple in  $U$ . In this way, all tuples from  $S$  are removed after looping over it the

<sup>2</sup>One may also recognize this as indexed nested loop join, which is equivalent [27].

<sup>3</sup>Backjumping is a concept in backtracking search algorithms; we use the term informally to mean the interruption of a nested loop iteration to jump back to an outer loop, while referring to the original TreeTracker algorithm [15] for a precise definition.

first time. Then, on all subsequent iterations of the loop over  $R$ , the probe into  $S$  fails immediately. Overall, the algorithm finishes in  $O(N)$  time.

In general, TTJ runs in linear time in the size of the input and output for full acyclic joins. But the algorithm is not limited to acyclic queries: given the same query plan for *any* query, cyclic or acyclic, TTJ is guaranteed to match the performance of binary hash join, when measuring query performance by counting the number of hash probes. In particular, when no probe fails TTJ behaves identically to binary join. This is in contrast to YA which always carries the overhead of semijoin reduction, even if the reduction does not remove any tuple.

To address cyclic queries further we introduce a new method to break down a cyclic query into acyclic parts called *tree convolution*. We use this method to analyze the run time of TTJ. A special kind of tree convolution, called *rooted convolution*, eliminates materialization of intermediates during query processing.

In summary, our contributions include:

- Propose TTJ, a new join algorithm that runs in time  $O(|IN| + |OUT|)$  on full acyclic queries.
- Prove that TTJ matches or outperforms binary join given the same query plan, on both acyclic and cyclic queries.
- Introduce *tree convolution*, a new method to break down cyclic queries into acyclic parts, and use it to analyze the run time of TTJ on cyclic queries.
- Improve the performance of TTJ with further optimizations.
- Conduct experiments to evaluate the efficiency of TTJ on acyclic queries.

## 2 Related Work

The observation that only one semijoin pass is necessary in YA has been a folklore in the database community, with an early appearance in the theoretical work of Bagan, Durand, and Grandjean [3]. Their paper studies the problem of enumerating conjunctive query results with constant delay, but without considering practical efficiency. Recent systems implementing such enumeration algorithms take advantage of the same insight [31, 32]. Compared to this improved version of YA, TTJ has the guarantee of matching the performance of HJ given any query plan, and is often faster in practice, as we will show in Section 6.

Since as early as the 80's, researchers have investigated different ways to implement sideway information passing [6, 14, 19, 26] to make database systems more robust. More recently, bitmap filters have been used for such purposes in Microsoft SQLSever [9, 17]. Compared to these approaches, TTJ requires no additional filter data structures and directly probes into the same hash tables used by binary hash join. This also allows us to provide the strong *zero-overhead* guarantee, where we prove that TTJ incurs no more hash table operations than binary hash join.

Researchers have also explored ways to integrate elements of YA into existing systems. Zhu et al. [41] propose *lookahead information passing*, using bloom filters to implement semijoins over star schemas. Birler, Kemper, and Neumann [5] decompose every join operator into a *lookup* and an *expand*, and prove that certain lookup-and-expand (L&E) plans are guaranteed to run in linear time for acyclic queries. Bekkers et al. [4] implement L&E plans in a vectorized query engine, while proving that their approach is guaranteed to match binary hash join for a class of *well-behaved* query plans. The theoretical guarantees of TTJ is complementary to these approaches: while TTJ guarantees to match binary hash join for left-deep plans, the well-behaved class defined by Bekkers et al. [4] essentially contains right-deep plans with a slight generalization.

Two very recent papers proposing practical instantiations of YA are due to Wang et al. [34] and Zhao et al. [40]. The *Yannakakis+* algorithm by Wang et al. [34] is based on query rewriting, which makes it easily applicable to any underlying database system. The *Robust Predicate Transfer*

algorithm by Zhao et al. [40] is a practical implementation of YA using bloom filters, which reduces the overhead of the semijoin reduction passes while allowing for a modular implementation in an existing system. Taking inspiration from these works, we design an optimized variant of TTJ, called  $TTJ^e$ , and compare it with these two algorithms. Each of these algorithms makes a different trade-off between ease of implementation and performance as we will discuss in detail in Section 6, and all three contribute valuable insights to bring the theoretical guarantees of YA to practice.

Wang et al. [35] and Idris et al. [13] propose optimal algorithms for *dynamic query processing*, where a system reacts to changes in the input by incrementally updating the output directly, without recomputing intermediates. Incidentally, these algorithms turn out to be instance-optimal when executed in *batch mode*, where each input relation is modeled as a batch of insertions provided upfront. TTJ can be understood as a *dual* to these algorithms. For example, the CROWN algorithm by Wang et al. [35] traverses the join tree bottom-up, performing semijoin reduction while building the hash table for each relation. In contrast, TTJ piggybacks semijoins during the *probe phase* after all base table hash tables are built. This allows building the hash tables in any order, in parallel, or even pre-built as indices. Furthermore, TTJ provably incurs no more hash table operations as compared to binary hash join, while bottom-up algorithms may require up to 2x hash table probes when the input database is already free of dangling tuples (in which case the semijoins are futile).

Going beyond acyclic queries, the standard way to handle cyclic queries is to break up the query with (hyper-)tree decomposition [10]. Such decomposition results in smaller cyclic subqueries connected by an acyclic “skeleton”. Each cyclic subquery can then be computed with worst-case optimal join algorithms [22, 33]. With the result of each subquery materialized, the final output can then be computed with YA. As we will show in Section 7, TTJ can support cyclic queries with only a few modifications. Compared to the tree decomposition approach, TTJ does not require materializing intermediate results, thus requiring only constant space in addition to the linear space required to store and index the input relations. While the worst-case time complexity of TTJ does not match that obtained by tree decompositions, the advantage of each approach depends on the data.

Wang et al. [37] proposed Free Join as an algorithm combining the strengths of worst-case optimal join and traditional hash join. The  $TTJ^e$  variant of our algorithm is directly inspired by Free Join, where we eagerly pull up hash table probes to avoid large intermediates. Our algorithms have stronger theoretical guarantees as compared to Free Join. Both TTJ and  $TTJ^e$  are instance-optimal for acyclic queries, while Free Join is not – the deletion step is crucial for instance-optimality.

As the name suggests, TreeTracker Join is a direct decendent of the TreeTracker algorithm [15] from Constraint Satisfaction. The TreeTracker CSP algorithm resolved Dechter’s conjecture [7] that there existed an optimal algorithm for acyclic CSPs free of any preprocessing. The connection between query answering and constraint satisfaction is a recurring theme in the literature to the extent that an expression emerged, the problems are two sides of the same coin [16, 20]. There are substantive differences that make TreeTracker and TTJ different. First the constraint satisfaction problem concerns the existence of a non-empty model for a large logical formula. Thus, constraint satisfaction algorithms including TreeTracker stop execution and return TRUE upon identifying what in a relational query would be just one row of the result. In contrast, TTJ produces all tuples in the query output. Second, the TreeTracker algorithm does not make use of hash tables, and is thus structured like a nested loop join rather than a hash join. This is because unlike the study of queries in databases, constraint satisfaction rarely specializes the problem to only equality predicates. Combining these two differences TreeTracker incorporates ad-hoc data structures, where TTJ employs recognized indices commonly used in databases. These difference clearly manifest in the respective complexity analyses. The complexity of the best variation of the TreeTracker algorithm

is polynomial in the input size and does not consider the output size. We prove below TTJ runs in linear time in the total size of the input and output.

### 3 Preliminaries

In this section, we present the foundational concepts concerning acyclic join queries and the specific definitions adopted in this paper.

#### 3.1 Join Queries and Acyclicity

We consider natural join queries, also known as *full conjunctive queries*, of the form:

$$Q(\mathbf{x}) = R_1(\mathbf{x}_1) \bowtie R_2(\mathbf{x}_2) \bowtie \cdots \bowtie R_n(\mathbf{x}_n) \quad (1)$$

where each  $R_i$  is a relation name, each  $\mathbf{x}_i$  (and  $\mathbf{x}$ ) a tuple of distinct variables, and every  $x \in \mathbf{x}_i$  also appears in  $\mathbf{x}$ . We call each  $R_i(\mathbf{x}_i)$  an *atom*, and  $\mathbf{x}_i$  the *schema* of  $R_i$ , denoted as  $\Sigma(R_i)$ . We extend the notion of schema to tuples in the standard way and write  $\Sigma(t)$  for the schema of  $t$ . The query computes the set<sup>4</sup>  $Q = \{\mathbf{x} \mid \bigwedge_{i \in [n]} \mathbf{x}_i \in R_i\}$ . We sometimes write  $Q$  and not  $Q(\mathbf{x})$  to reduce clutter, and identify  $Q$  with its set of relations. For example,  $Q - \{R_i\}$  denotes the query  $Q$  with  $R_i$  removed.

*Definition 3.1 (Join Tree).* A *join tree* for a query  $Q$  is a tree where each node is an atom in  $Q$ , with all atoms in  $Q$  appearing in the tree, and for every variable  $x$ , the nodes containing  $x$  form a connected subtree.

A query  $Q$  is *acyclic* (more specifically  $\alpha$ -acyclic) if there exists a *join tree* for  $Q$ .

For clarity we rewrite the query in example 1.1 and detail one of its join trees:

$$Q_1 = R(i, x) \bowtie S(x, y, j) \bowtie T(y, k) \bowtie U(y, l) \quad (2)$$

One join tree has  $R(i, x)$  at the root,  $S(x, y, j)$  as its child, and  $T(y, k)$  and  $U(y, l)$  as children of  $S$ . We encourage the reader to draw a picture of this join tree for reference. One can construct a join tree for any acyclic query with the GYO algorithm [12, 39], which works by finding a sequence of *ears*. To define ear, we first define a *key schema*:

*Definition 3.2 (Key Schema).* For a query  $Q$  of the form (1), the *key schema* of an atom  $R_i(\mathbf{x}_i)$  in  $Q$ , denoted as  $\text{keys}(Q, R_i)$ , is the set of variables shared between  $R_i(\mathbf{x}_i)$  with the other atoms in  $Q$ ; i.e.,  $\text{keys}(Q, R_i) = \mathbf{x}_i \cap \bigcup_{j \in [n] \wedge j \neq i} \mathbf{x}_j$ .

Intuitively,  $\text{keys}(Q, R_i)$  form the keys of  $R_i$ 's hash table, if we compute  $(Q - \{R_i\}) \bowtie R_i$  using binary hash join.

*Definition 3.3 (Ear).* Given a query  $Q$  of the form (1), an atom  $R_i(\mathbf{x}_i)$  is an *ear* if it satisfies the property  $\exists j \neq i : \mathbf{x}_j \supseteq \text{keys}(Q, R_i)$ . In words, there is another atom  $R_j(\mathbf{x}_j)$  that contains all the variables in  $R_i$ 's key schema. We call such an  $R_j$  a *parent* of  $R_i$ .

The parent concept is central to the TTJ algorithm. The parent's schema include all of its children's keys. When a hash lookup fails at a child, TTJ will backjump to the parent. Figure 2b shows an algorithm to find the first parent of an ear in  $Q$ , where  $Q$  is represented as a list of atoms.

The GYO algorithm for constructing join trees is shown in Figure 2a: we start with a forest where each atom makes up its own tree, then for every ear, we attach it to its parent and remove that ear from the query. Note that it is possible for the algorithm to produce a forest of disjoint trees when the query contains Cartesian products. For simplicity, we will ignore such cases.

<sup>4</sup>For clarity we assume set semantics. No change is needed for TTJ to support bag semantics

```

1 def GYO(Q):
2   forest = { tree(R) for R in Q }
3   while not Q.is_empty():
4     R = find-ear(Q)
5     P = parent(Q, R)
6     forest.set_parent(R, P)
7     Q.remove(R)
8   return forest

```

(a) The GYO algorithm.

```

1 def parent(R, Q):
2   if Q.is_empty(): return None
3   keys =  $\Sigma(R) \cap \bigcup_{S \in Q - \{R\}} \Sigma(S)$ 
4   for S in Q - {R}:
5     if keys  $\subseteq \Sigma(S)$ :
6       return S
7   # we did not find a valid parent
8   return None

```

(b) Find a parent of  $R$  in  $Q$  if one exists.

Fig. 2. GYO reduction and parent computation.

```

1 def join(t, plan, i):
2   if i > plan.len(): print(t)
3   else:
4     R = plan[i]; k =  $\pi_{\text{keys}}(\text{plan}[1..i], R)(t)$ 
5     for r in R[k]:
6       join(t++r, plan, i+1)

```

(a) Pipelined left-deep binary hash join

```

1 def YA(Q, order):
2   for R in order: # semijoins reduction
3     P = parent(Q, R); Q.remove(R)
4     if P is not None: P = P  $\bowtie$  R
5   # compute the output with hash join
6   return join((), reverse(order), 1)

```

(b) Yannakakis's algorithm

Fig. 3. Binary hash join and Yannakakis's algorithm. The `plan` array is 1-indexed.

*Definition 3.4 (GYO reduction order).* Given a query  $Q$  of the form (1), a *GYO reduction order* for a query  $Q$  is a sequence  $[R_{p_1}, R_{p_2}, \dots, R_{p_n}]$  that is a permutation of  $[R_1, R_2, \dots, R_n]$ , such that for every  $i < n$ , the atom  $R_{p_i}$  is an ear in the (sub)query  $R_{p_i} \bowtie \dots \bowtie R_{p_n}$ .

Equivalently, it is the same order of atoms as visited by the GYO algorithm. The reader can verify  $[U, T, S, R]$  is a GYO reduction order for  $Q_1$ . The existence of a GYO reduction order and the existence of a join tree are equivalent.

**THEOREM 3.5** ([12, 39]). *A query  $Q$  has a join tree (i.e.,  $Q$  is  $\alpha$ -acyclic) if and only if it has a GYO reduction order.*

### 3.2 Binary Hash Join

In this paper we focus on hash-based join algorithms. For theoretical analyses we focus on left-deep linear plans; for practical implementation we follow the standard practice and decompose each bushy plan into a sequence of left-deep linear plans, materializing each intermediate result.

*Definition 3.6 (Query Plan).* A (left-deep linear) query plan for a query  $Q$  of the form (1) is a sequence  $[R_{p_1}, R_{p_2}, \dots, R_{p_n}]$  that is a permutation of  $Q$ 's relations  $[R_1, R_2, \dots, R_n]$ .

For consistency we adopt 1-based indexing for query plans, so the first relation in the plan is stored at  $i = 1$ . An example query plan for  $Q_1$  in Equation (2) is  $[R, S, T, U]$ . One may notice similarities between a GYO reduction order and a query plan. The reason for this will become clear.

We follow the push-based model [21] and specialize the binary hash join algorithm for pipelined left-deep plans as shown in Figure 3a. We write  $\pi_s(t)$  to project the tuple  $t$  onto the schema  $s$ , and  $t \bowtie r$  to concatenate the tuples  $t$  and  $r$  while resolving the schema appropriately. Execution begins by passing to join the empty tuple  $t = ()$ , a query plan, and  $i = 1$ . Although we do not need to build

a hash table for the left-most relation (the first relation in the plan), for simplicity we assume that there is a (degenerate) hash table mapping the empty tuple  $()$  to the entire left-most relation. The algorithm starts by checking if the plan has been exhausted and if so, output the tuple  $t$ . Otherwise, we retrieve the  $i$ -th relation  $R_{p_i}$  from the plan, and lookup from  $R_{p_i}$  the matching tuples that join with  $t$ . For each match, we concatenate it with  $t$  and recursively call `join`.

It may be helpful to unroll the recursion over a query plan, and we encourage the reader to do so for  $Q_1$  in (2) with the plan  $[R, S, T, U]$ . This will generate the same code as in Figure 1a.

### 3.3 Yannakakis's Algorithm

Yannakakis's original algorithm [38] makes two preprocessing passes over the input relations. A third pass computes the joins yielding the final output. Bagan, Durand, and Gandjean [2] improved the original algorithm by eliminating the second preprocessing pass. For brevity we only describe the latter algorithm. Following common usage, hereafter, we will refer to the improved version as Yannakakis's algorithm (YA).

Shown in Figure 3b is, given a GYO reduction order, the relations are preprocessed using semijoins, then the output is computed with standard hash join. Equivalently, the semijoin preprocessing step can be performed by traversing a join tree bottom-up, and the output computed with hash join by traversing the tree top-down.

*Example 3.7.* Given the query  $Q_1$  in (2) and the GYO reduction order  $[U, T, S, R]$ , YA first performs the series of semijoins,  $S' = S \ltimes U$ ,  $S'' = S' \ltimes T$ , and  $R' = R \ltimes S''$ , then computes the output with the plan  $[R', S'', T, U]$ . The reader may refer to the join tree of  $Q_1$  and confirm we are traversing the tree bottom-up then top-down.

## 4 TreeTracker Join

The TreeTracker Join algorithm is shown in Figure 4a. The algorithm follows the same structure as binary hash join. The difference starts on line 5 right before the hash lookup  $R[k]$ . If this lookup fails (i.e., it finds no match), and if  $R$  has a parent  $P$  that appears before  $R$  in the plan, then TTJ backjumps to the `for`-loop at  $P$ 's recursive level, by returning  $P$  (line 7). This is similar to throwing an exception which is “caught” at the loop level of  $P$ , as we will explain on line 10. Otherwise, if the lookup  $R[k]$  succeeds, the algorithm iterates over each matching tuple  $r$  and calls itself recursively (line 9). This recursive call has three possible results. A result containing a relation (line 10) signifies a backjump has occurred, with that relation as the backjumping point. If the backjumping point is the same as the current relation  $R$ , then the tuple  $r$  is deleted from  $R$  (line 11). If the backjumping point is different from  $R$ , then the backjump continues by returning `result` which interrupts the current loop. Finally, if the recursive call (implicitly) returns `None`, the algorithm continues to the next loop iteration.

*Example 4.1.* It can be helpful to unroll the recursive algorithm over a query plan. Given  $Q_1$  in (2) and the plan  $[R, S, T, U]$ , Figure 4b shows the execution of TTJ. To make the code more intuitive, we replace `return` statements with exception handling to simulate backjumping. We gray out dead code and no-ops:

- Line 1 is unreachable because  $R[()]$  is always the entire relation  $R$ , and  $R$  has no parent.
- Line 3 (and 13) is a no-op, because it would just backjump to the immediately enclosing loop, and removing a tuple from  $R$  is useless because  $R$  is at the outermost loop<sup>5</sup>.

<sup>5</sup>In Section 5 we will introduce an additional optimization that makes “removing” from the outermost relation meaningful.

```

1  def ttj(t, plan, i):
2      if i > plan.len(): print(t)
3      else:
4          R = plan[i]; k =  $\pi_{\text{keys}(\text{plan}[1..i], R)}(t)$ 
5          P = parent(plan[1..i], R)
6          if R[k] is None & P is not None:
7              return P # backjump to P
8          for r in R[k]:
9              result = ttj(t+r, plan, i+1)
10             if result == R: # catch backjump
11                 R[k].delete(r)
12             elif: result is not None:
13                 return result # continue backjump

```

(a) TreeTracker join

### (b) Execution of TTJ for Example 1.1

Fig. 4. The TreeTracker algorithm and an example execution.

- Technically the if-statement on line 5 is useful even though it only backjumps one level, because the backjump would remove a tuple from  $S$  when caught (line 12). However for the input data in Example 1.1 we do not need this, and we gray it out to reduce clutter.
- Finally, the innermost two try-catch pairs are unreachable, because  $U$  and  $T$  have no children.

At this point, the remaining code in black is essentially the same as the code in Figure 1c. As a side note, a sufficiently smart compiler with partial evaluation or just-in-time compilation could remove the dead code and no-ops as we have done above.

## 4.1 Correctness and Asymptotic Complexity

The correctness proof starts with an observation on the relationship between different calls to `ttj`:

PROPOSITION 4.2. If  $\text{ttj}(t_j, p, j)$  recursively calls  $\text{ttj}(t_i, p, i)$ , then  $t_j \subseteq t_i$ .

PROOF. The proposition follows from the definition of the algorithm, where the  $t_i$  argument to the nested call is constructed by appending tuples to  $t_j$ .  $\square$

TTJ differs from binary join only upon a lookup failure. In that case it backjumps to the parent of the relation that caused the failure, and deletes the tuple that caused the failure. Therefore, TTJ is correct as long as it never deletes or “backjumps over” any tuple that should be in the output. We first prove that a deleted tuple can never contribute to any output. In the following we write  $\pi_R(t)$  for the projection of  $t$  onto the schema of  $R$ .

LEMMA 4.3. Suppose a tuple  $r_j$  is deleted from  $R_j$  during the execution of TTJ for a query  $Q$  using plan  $p$ . Then  $\forall t_{out} \in Q : \pi_{R_j}(t_{out}) \neq r_j$ .

PROOF. Let  $p$  be  $[R_1, \dots, R_n]$ , and  $t_j$  be the value of the argument  $t$  in scope at the time of the deletion. Because  $r_j$  is deleted from  $R_j$ , there must be a failed lookup  $R_i[k_i]$  recursively nested within the call to  $\text{ttj}(t_j + r_j, p, j + 1)$ , and  $R_j$  is the parent of  $R_i$ . Let  $K_i = \text{keys}(p[1, \dots, i], R_i)$ , and let  $t_i$  be the value of  $t$  at the time of the lookup failure. Then  $t_j + r_j \subseteq t_i$  by Proposition 4.2. By definition of parent,  $K_i \subseteq \Sigma(R_j) \subseteq \Sigma(t_j + r_j) \subseteq \Sigma(t_i)$ , so  $k_i = \pi_{K_i}(t_i) = \pi_{K_i}(t_j + r_j) = \pi_{K_i}(r_j)$ . However, since the lookup failure implies no tuple in  $R_i$  contains  $k_i$ , any output tuple  $t_{\text{out}}$  cannot

contain  $k_i$  either, i.e.,  $\forall t_{\text{out}} \in Q : \pi_{K_i}(t_{\text{out}}) \neq k_i$ . Therefore,  $\forall t_{\text{out}} \in Q : \pi_{K_i}(t_{\text{out}}) \neq \pi_{K_i}(r_j)$  which implies  $\forall t_{\text{out}} \in Q : \pi_{R_j}(t_{\text{out}}) \neq \pi_{R_j}(r_j)$ .  $\square$

Next, we show TTJ never backjumps over any tuple that contributes to the output. Given a plan  $p = [R_1, \dots, R_n]$ , denote by  $\pi_{[i]}(t)$  the projection of  $t$  onto  $\bigcup_{j \in [i]} \Sigma(R_j)$ .

**LEMMA 4.4.** *For any tuple  $t_{\text{out}} \in Q$ , plan  $p$  for  $Q$ , and  $1 \leq i \leq |p|$ ,  $\text{ttj}(\pi_{[i-1]}(t_{\text{out}}), p, i)$  recursively calls  $\text{ttj}(\pi_{[i]}(t_{\text{out}}), p, i + 1)$ .*

**PROOF.** Consider a lookup  $R[k]$  that is recursively nested within the call to  $\text{ttj}(\pi_{[i-1]}(t_{\text{out}}), p, i)$  where  $R$  has a parent  $R_j$  with  $j \in [i-1]$ . Then  $k \subseteq \pi_{R_j}(t_{\text{out}}) \subseteq t_{\text{out}}$ , and because  $\Sigma(k) \subseteq \Sigma(R)$ , we have  $k \subseteq \pi_R(t_{\text{out}}) \in R$ . This means the lookup  $R[k]$  will not fail. This holds for all such  $R$ , so the algorithm never backjumps from within the call  $\text{ttj}(\pi_{[i-1]}(t_{\text{out}}), p, i)$  to any  $R_j$  for  $j \in [i-1]$ . The algorithm may still backjump to  $R_i$ , but by Lemma 4.3,  $\pi_{R_i}(t_{\text{out}})$  is never deleted from  $R_i$ , and therefore the algorithm will recursively call  $\text{ttj}(\pi_{[i-1]}(t_{\text{out}}) \uparrow\!\!\uparrow \pi_{R_i}(t_{\text{out}}), p, i + 1)$  which is the same as  $\text{ttj}(\pi_{[i]}(t_{\text{out}}), p, i + 1)$ .  $\square$

We arrive at the correctness of TTJ by applying Lemma 4.4 inductively over the query plan.

**THEOREM 4.5.** *Given any plan  $p$  for  $Q$ ,  $\text{ttj}(\emptyset, p, 1)$  computes  $Q$ .*

**PROOF.** We prove the correctness of TTJ in two directions: first, any tuple produced by TTJ should be in the output; second, TTJ produces all tuples that should be in the output. The first direction is straightforward, as any tuple produced by TTJ is also produced by binary hash join. We prove the second direction by induction over the argument  $i$ , with the following inductive hypothesis:  $\text{ttj}(\pi_{[i-1]}(t_{\text{out}}), p, i)$  will be invoked for all  $t_{\text{out}} \in Q$  and  $1 \leq i \leq |p|$ . The base case when  $i = 1$  holds because we start the execution of TTJ by calling  $\text{ttj}(\emptyset, p, 1)$ . For the inductive step, assume  $\text{ttj}(\pi_{[i-1]}(t_{\text{out}}), p, i)$  is invoked, then applying Lemma 4.4 shows  $\text{ttj}(\pi_{[i]}(t_{\text{out}}), p, i + 1)$  will also be invoked. Therefore,  $\text{ttj}(t_{\text{out}}, p, |p| + 1)$  will be invoked for all  $t_{\text{out}} \in Q$ , which produces all tuples that should be in the output.  $\square$

Next, we prove TTJ runs in linear time in the size of the input and output, for full acyclic queries. We first introduce a condition on the query plan that is necessary for the linear time complexity:

**LEMMA 4.6.** *Given a query  $Q$  and a plan  $p = [R_1, \dots, R_n]$  for  $Q$ , parent returns None only for  $R_1$  during the execution of TTJ, if  $p$  is the reverse of a GYO reduction order of  $Q$ .*

**PROOF.** If  $[R_n, \dots, R_1]$  is a GYO reduction order, then there is a join tree with  $R_1$  as root, and every non-root atom has a parent.  $\square$

TTJ is guaranteed to run in linear time given such a plan:

**THEOREM 4.7.** *Fix a query  $Q$  and a plan  $p$ . If  $p$  is the reverse of a GYO reduction order for  $Q$ , then  $\text{ttj}(\emptyset, p, 1)$  computes  $Q$  in time  $O(|Q| + \sum_i |R_i|)$ .*

**PROOF.** We first note that in Figure 4a,  $\text{ttj}$  does constant work outside of the loops; each iteration of the loop also does constant work and recursively calls  $\text{ttj}$ , so each call to  $\text{ttj}$  accounts for constant work, therefore the total run time is linear in the number of calls to  $\text{ttj}$ . All we need to show now is that there are a linear number of calls to  $\text{ttj}$ .

Because  $p$  is the reverse of a GYO reduction order for  $Q$ , the following holds from Lemma 4.6: except for the one call to  $\text{ttj}$  on the root relation (when  $i = 1$ ), every call to  $\text{ttj}$  has 3 possible outcomes: (1) It outputs a tuple. (2) It backjumps and deletes a tuple from an input relation. (3) It recursively calls  $\text{ttj}$ . Because the query plan has constant length, there can be at most a constant number of recursive calls to  $\text{ttj}$  (case 3) until we reach cases 1 or 2. Therefore there are at most  $O(|Q| + \sum_i |R_i|)$  calls to  $\text{ttj}$ , and the algorithm runs in that time.  $\square$

By Theorem 3.5 every  $\alpha$ -acyclic query can be GYO-reduced, therefore  $\text{ttj}$  runs in linear time:

**COROLLARY 4.8.** *For any  $\alpha$ -acyclic query  $Q$ , there is a plan  $p$  such that  $\text{ttj}(\text{()}, p, 1)$  computes  $Q$  in time  $O(|Q| + \sum_i |R_i|)$ .*

## 4.2 Comparison with Binary Join and YA

We now prove our claim that, for any given query plan, TTJ always matches or outperforms binary hash join. Because TTJ and hash join build the exact same set of hash tables, they share the same cost for hash building. We therefore focus on the cost of hash lookups which accounts for the majority of the remaining cost for both algorithms. The following proofs take advantage of set semantics, but it is easy to extend the reasoning for bag semantics, as we can convert a bag into a set by appending a unique labeled null value to each tuple. We start with the following observation to relate the run time of hash join and TTJ to the set of arguments they are invoked with:

**LEMMA 4.9.** *Both hash join and TTJ, as defined in Figure 3a and Figure 4a, are invoked once for each distinct combination of the arguments  $(t, p, i)$ .*

**PROOF.** We prove by induction over the argument  $i$ . In the base case when  $i = 1$ , both algorithms are invoked once with  $t = \text{()}$ ,  $i = 1$ . For the inductive step, first consider the hash join algorithm. For every distinct  $t$ ,  $\text{join}(t, \text{plan}, i)$  recursively calls  $\text{join}(t+r, \text{plan}, i+1)$  for every  $r \in R_i[k]$ . Since  $R_i$  is a set, each  $r$  is distinct, so each  $t+r$  is also distinct. The same reasoning also applies to TTJ, as the algorithm will call itself only for a subset of the tuples in  $R_i[k]$ .  $\square$

In other words, the number of calls to each algorithm is the same as the number of distinct arguments they are invoked with. We can now compare the algorithms, by bounding the number of calls to TTJ by that of binary join.

**THEOREM 4.10.** *Given a query  $Q$  and a plan  $p$  for  $Q$ , computing  $Q$  with TTJ using  $p$  makes at most as many hash lookups as computing  $Q$  with binary join using  $p$ .*

**PROOF.** For clarity we have repeated the lookup  $R[k]$  three times in Figure 4a, but we really only need to look up once and save the result to a local variable for reuse. Specifically, a pointer to  $R[k]$  on line 6 can be used for the nullness check on the same line, the loop on line 8, as well as the deletion<sup>6</sup> on line 11. This way, every call to  $\text{ttj}$  makes exactly one hash lookup. Since the binary join algorithm in Figure 3a also makes exactly one hash lookup per call, it is sufficient to bound the number of calls to  $\text{ttj}$  by that of binary join. By Lemma 4.9, it is sufficient to show the distinct arguments TTJ is invoked on is a subset of that for binary join. We prove this by induction over the argument  $i$ . When  $i = 1$ , both TTJ and binary join are invoked with  $t = \text{()}$  and  $i = 1$ . For the inductive step,  $\text{ttj}(t, p, i)$  recursively calls  $\text{ttj}(t+r, p, i+1)$  only if  $r \in R_i[k]$ , which implies  $\text{join}(t, p, i)$  will also call  $\text{join}(t+r, p, i+1)$  in binary join. Therefore, every call to TTJ is accounted for with a call to binary join.  $\square$

Another cost in query execution comes from accessing the matching tuples after a successful lookup, and one can prove that TTJ accesses no more tuples than binary join, following the same reasoning as above. Although backjumping and tuple deletion in TTJ may in principle carry an overhead, we will show in Section 6 that such an overhead is negligible as compared to the cost of hash lookups. Finally, we note the above proof does not assume an acyclic query. Section 7 analyzes the run time of TTJ on cyclic queries.

<sup>6</sup>Although the deletion occurs after a recursive function call, the recursion has constant depth, so the pointer dereference has good temporal locality and is likely cheap.

While we guarantee TTJ to always match binary join, we cannot make the same strong claim for YA. We will see in Section 6 that YA performs better than TTJ on some queries. Here we analyze a few extreme cases for some intuition of how TTJ compares to YA:

*Example 4.11.* Consider a query where every tuple successfully joins, i.e., no lookup fails. In this case binary join and TTJ behaves identically. However, YA spends additional time futilely computing semijoins (without removing any tuple), before following the same execution as binary join and TTJ to produce the output.

*Example 4.12.* The other extreme case is when a query has no output, and YA immediately detects this and stops. In fact Example 1.1 is such a query: all YA needs to do is the semijoin  $T \bowtie U$ , where it builds a (tiny) hash table for  $U$  and iterate over  $T$  once to detect nothing joins. In contrast, although TTJ also runs in linear time, it must build the hash table for all of  $S$ ,  $T$  and  $U$ .

### 4.3 Cartesian Products and Non-equality Joins

While we have focused on equality joins to simplify presentation, TTJ supports both cartesian products and non-equality joins. If the query contains a Cartesian product, TTJ would simply never backjump across the product operator, and the guarantees of instance-optimality and zero-overhead continue to hold. The simplest way to handle non-equality joins is to apply non-equality predicates after all equality joins, while range predicates can be seen as a generalization of hash lookup: if a range predicate matches no results in a B+ tree index, we can safely backjump and delete just like TTJ does for hash lookup failures. A recent work by Wang and Yi [36] presents optimal algorithms for conjunctive queries with comparisons. Extending TTJ to achieve similar optimality for such queries is an interesting direction for future work.

## 5 Practical Considerations

Although TTJ has strong guarantees, its reliance on backjumping and deletion stands in the way of an efficient implementation. Backjumping complicates the control flow and therefore hinders compiler optimizations. Similarly, interleaving deletion with lookups also leads to more intricate code (e.g., requiring a `while` loop instead of a simpler `for each` loop), and is incompatible with vectorization and parallelization. We address these issues in this section with an alternative formulation of TTJ that trades off some of the theoretical guarantees for practical performance. Our approach is inspired by earlier work using bloom filters and bitmap filters to implement various forms of sideway information passing [6, 14, 19, 26]. However, unlike these prior work, we require no additional filter data structures and directly probe into the same hash tables as used by binary hash join. We discuss an additional optimization called no-good list in Appendix A.

### 5.1 Eliminating Backjumping

The main goal of backjumping in TTJ is to avoid iterating over match results before performing all applicable lookups. By jumping back upon a lookup failure, TTJ guarantees the same failure will not be repeated. Inspired by past work on pre-filtering using bitmap filters or bloom filters, we replace TTJ’s backjumping with *early probing*: for every relation  $R$ , we immediately probe into all children of  $R$  before entering any subsequent loops. Unlike prior work on pre-filtering, our approach does not require any additional data structures like bitmaps or bloom filters. Furthermore, early probing still guarantees the resulting algorithm is zero-overhead to a certain extent, which we will discuss in detail below. We note a similar technique was used in the Free Join algorithm by Wang et al. [37] to implement worst-case optimal join for cyclic queries. However, Free Join is not instance-optimal for acyclic queries.

Figure 5a shows the pseudocode for TTJ with early probing, denoted as  $\text{TTJ}^e$ . For each relation  $R$  in the plan, we find all children relations  $C$  such that  $R$  is a parent of  $C$ . Then for each tuple  $r$  in  $R$ , we probe into each child relation  $C$ . Upon any probe failure, we delete  $r$  from  $R$  and skip to the next tuple. If all probes succeed, we recursively call  $\text{TTJ}^e$  after replacing each  $C$  with the lookup result  $C[k]$ . Figure 5b shows an example execution of  $\text{TTJ}^e$  for the query in Example 1.1. The root relation  $R$  has only one child  $S$ , which is probed into at the top loop level. If the probe fails, the tuple  $(i, x)$  is deleted from  $R$ . Otherwise, we proceed to iterate over the matching tuples  $(y, j)$  in  $S[x]$ .  $S$  has two children  $T$  and  $U$ , and we probe into them at once. If either probe fails, the tuple  $(y, j)$  is deleted from  $S[x]$ . Otherwise, we iterate over the matching tuples and output the final result. Note that the execution example no longer contains any non-local loop exits. A vectorized implementation may therefore perform lookups in batch; the simplified control flow also paves the way for parallelization.

It is easy to see  $\text{TTJ}^e$  runs in linear time for acyclic queries, following the same reasoning in the proof of Theorem 4.7:

**THEOREM 5.1.** *Fix a full acyclic query  $Q$  and a plan  $p$ . If  $p$  is the reverse of a GYO reduction order for  $Q$ , then  $\text{ttj-e}(\cdot, p, 1)$  computes  $Q$  in time  $O(|Q| + \sum_i |R_i|)$ .*

**PROOF SKETCH.** The body of the outer loop in Figure 5a does constant work per iteration (when we regard the size of query as constant), and each iteration either recursively calls  $\text{TTJ}^e$  or deletes a tuple. Since the algorithm will produce an output after a constant number of recursive calls, the entire algorithm must run in time  $O(|\text{IN}| + |\text{OUT}|)$ .  $\square$

Compared to TTJ which guarantees to incur no more hash probes than HJ on *any* query plan,  $\text{TTJ}^e$ 's guarantee is less general. Intuitively, the early probing in  $\text{TTJ}^e$  essentially reorders joins. Although this reordering guarantees asymptotic optimality, it can introduce constant-factor overhead on the specific input. Nevertheless,  $\text{TTJ}^e$  still guarantees zero-overhead for query plans that are “compatible” with a join tree (the difference with Theorem 4.10 is highlighted in bold):

**THEOREM 5.2.** *Given a query  $Q$  and a plan  $p$  for  $Q$  such that **for every  $R \in Q$  and every child  $C$  of  $R$ , any relation  $S$  between  $R$  and  $C$  in  $p$  is also a child of  $R$** , computing  $Q$  with  $\text{TTJ}^e$  using  $p$  makes at most as many hash lookups as computing  $Q$  with binary join using  $p$ .*

**PROOF SKETCH.** If the query plan satisfies the above conditions,  $\text{TTJ}^e$  performs hash lookups in the same order as binary hash join, but avoids fruitless iteration over matched results when a lookup fails.  $\square$

## 5.2 Avoiding Deletion

While  $\text{TTJ}^e$  completely eliminates backjumping, deletion is still necessary to guarantee linear time complexity for the general case. Nevertheless, we can avoid deletion as much as possible by exploiting the structure of the query as well as primary/foreign key constraints in the data. In this section we present several of this kind of *semantic optimizations*.

**Shallow Join Trees.** We have observed in Example 4.1 that deletion from the root relation of a join tree is unnecessary, because the root relation appears at the top loop level, and every tuple will only be considered once. Therefore, we can eliminate all deletions if the join tree has a depth of two, i.e., all non-root relations have the root as parent. This is the case for star-schema queries, where we may choose the fact table as root, and the dimension tables as children. In general, when generating query plans, we favor shallow join trees over deep ones. This is implemented with the maximum cardinality search algorithm by Tarjan and Yannakakis [29] which constructs a join tree

```

1  def ttj-e(t, plan, i):
2      if i > plan.len(): print(t)
3      else:
4          R = plan[i]
5          l1: for r in R:
6              new_plan = clone(plan)
7              for C in children(plan, R):
8                  k = πR∩C(r)
9                  if C[k] is None:
10                     R.delete(r); continue l1
11                  else: new_plan[C\k] = C[k]
12          ttj-e(t+r, new_plan, i+1)
13
14      for i,x in R:
15          if S[x] is None:
16              R.delete(i, x)
17          else:
18              for y,j in S[x]:
19                  if T[y] is None || U[y] is None:
20                      S[x].delete(y,j)
21                  else:
22                      for k in T[y]:
23                          for l in U[y]:
24                              output(x,y,i,j,k,l)
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100
101
102
103
104
105
106
107
108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215
216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269
270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323
324
325
326
327
328
329
330

```

(a) TTJ with early probing (TTJ<sup>ε</sup>)(b) Execution of TTJ<sup>ε</sup> for Example 1.1Fig. 5. TTJ with early probing (TTJ<sup>ε</sup>) and example execution.

root-to-leaf, while greedily attaching each child to the highest possible parent. After we generate the structure of the join tree, we order the children of each relation based on cardinality estimates.

*Primary/Foreign Key Constraints.* The justification for skipping deletion at the root is that the same tuple will never be considered again. We can extend this reasoning to non-root relations that join with the root via primary/foreign keys. Specifically, if a relation  $S$  joins with the root  $R$  with the condition  $S.f = R.p$  where  $p$  is a primary key of  $R$  and  $f$  is a foreign key of  $S$  referencing  $p$ , then every tuple in  $S$  will only be visited once, and it is therefore safe to skip deletion from  $S$ . Another way to exploit primary/foreign key constraints is to pre-compute intermediate results that are guaranteed to have linear size. For example, if the query plan contains  $S \bowtie T$  as a subplan, and if that join is on a primary/foreign key, then the output size will not exceed the total size of  $S$  and  $T$ . We can therefore safely materialize this join (using regular binary hash join), which in turn makes the join tree shallower and can further eliminate deletion.

## 6 Empirical Results

Since our primary contribution concerns the development of an algorithm that is both asymptotically optimal and is competitive in practice w.r.t. wall clock time, the primary goal of the empirical assessment is to compare the execution time of the algorithms in as controlled of an experiment as possible. All three algorithms, TTJ, binary hash join, HJ, and YA are implemented in the same Java query execution engine written from scratch. We are certain our algorithm execution measurements do not make calls to methods outside of our execution environment. Any data structure in our execution environment whose definition is impacted by the definition of a data structure outside of our Java execution environment is treated identically for all three algorithms. Where possible, code is reused across algorithm implementation. The source code of the implementation is available at <https://anonymous.4open.science/r/treetracker>.

Remaining aspects of query compilation and DBMS implementation are “borrowed” from other DBMS implementations. Query plans are an example of borrowing from other DBMS implementations. After loading a benchmark database instance and gathering catalog statistics left-deep linear query plans are determined by SQLite, and bushy plans by PostgreSQL. The SQL EXPLAIN command elicits the plans from the DBMSs. SQLite and PostgreSQL were chosen because of the topology of the plans their optimizers generate. The linear time guarantee only holds for left-deep

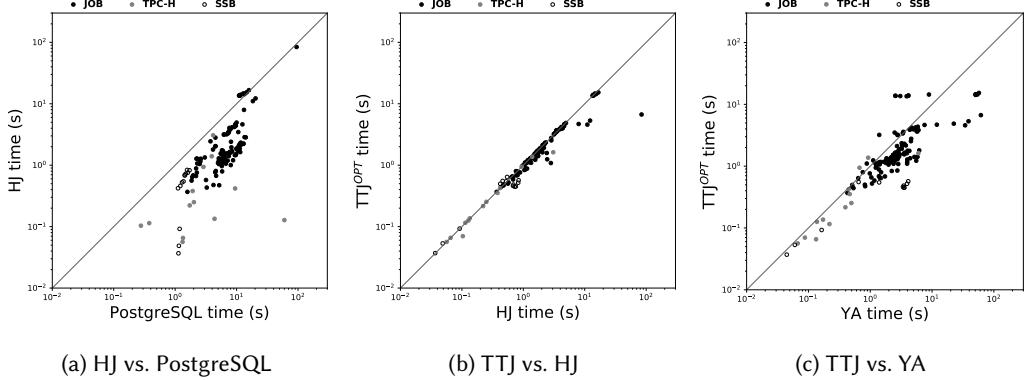


Fig. 6. Run time of TTJ, HJ, YA, and PostgreSQL on JOB, TPC-H, and SSB. Every data point corresponds to a query, whose  $x$ - and  $y$ -coordinates correspond to the run time of the algorithms under comparison.

linear plans that are consistent with a GYO reduction order of the query. All the left-deep plans produced by SQLite in our experiments are consistent with the GYO reduction requirement.

*Workload.* Our experiments encompass left-deep plans, left-deep plans with optimizations integrated into the TTJ algorithm, and bushy plans. Only the acyclic join queries in three benchmarks were evaluated, the Join Ordering Benchmark (JOB) [18], TPC-H [30] (scale factor = 1), and the Star Schema Benchmark (SSB) [24] (scale factor = 1). Also omitted were single-relation queries, and correlated subqueries. These criteria eliminated only 9 queries, all from TPC-H. Thus, the 113 JOB queries, the 13 SSB queries and 13 out of 22 TPC-H queries were assessed, for a total of 139 queries.

*Environment.* Experiments were conducted on a single logical core of an AMD Ryzen 9 5900X 12-Core Processor @ 3.7Hz CPU. The computer contained 64 GB of RAM, and a 1TB PCIe NVMe Gen3 M.2 2280 Internal SSD. All data structures are allocated from JVM heap which was set to 20 GB. Since execution was otherwise identical for all algorithms under test, no techniques to reduce the overhead of memory allocation or garbage collection were exploited. Measurements for each query and algorithm pair were orchestrated by JMH [1] configured for 5 warmup forks and 10 measurement forks. Each of those forks contains 3 warmup and 5 measurement iterations.

Direct measurements of PostgreSQL, which can be seen as the control group (not a baseline), for our implementations are the same as [18]. PostgreSQL measurements use an in-memory hash join, indices were dropped and single process execution specified. Thus, we configured PostgreSQL such that measurements were made as similar to our Java implementations as we could make possible. A timeout was set to 1 minute. Of all the queries only TPC-H Q20 exceeded the timeout.

## 6.1 Algorithm Comparison

Figure 6 illustrates our primary results. It contains 3 scatter plots that pairwise compare the execution time of 4 implementations for each query across the 3 benchmarks. First, Figure 6a compares the performance of PostgreSQL, using hash joins with our implementation of HJ. Inspection of the scatterplot shows that with few exceptions the execution time of the same query is less than an order of magnitude apart. Most points are below the diagonal indicating our implementation is faster than PostgreSQL. The shape of the cluster suggests a consistent range in the disparity of execution time.

Faster execution is not surprising. The results of a road race with PostgreSQL are not material to this paper. Our execution environment contains no elements of transaction system overhead or

Baseline	Benchmark	Average	Maximum	Mininimum
Hash Join	JOB	1.11 $\times$	12.6 $\times$ (16b)	0.9 $\times$ (11b)
	TPC-H	1.09 $\times$	1.9 $\times$ (Q9)	1 $\times$ (Q7)
	SSB	1.15 $\times$	1.7 $\times$ (Q2.2)	0.8 $\times$ (Q3.4)
Yannakakis's Algorithm	JOB	1.60 $\times$	9.2 $\times$ (16b)	0.2 $\times$ (6a)
	TPC-H	1.40 $\times$	3.7 $\times$ (Q9)	0.7 $\times$ (Q7)
	SSB	3.16 $\times$	7.9 $\times$ (Q2.2)	1 $\times$ (Q3.4)

Table 1. Speed-up of TreeTracker Join Relative to Hash Join and Yannakakis's Algorithm.

buffer and memory hierarchy management. PostgreSQL execution time was measured as a control. This first plot establishes that our Java implementation is within range of a commercially used RDBMS and the consistency in the difference of execution speed lends credibility that the empirical results from our execution environment will generalize to commercially deployed RDBMSs.

Figure 6b shows, on a per query basis, the relative speed of TTJ versus HJ. The visualization in Figure 6b reveals that TTJ is often faster than HJ, and for just a few queries the execution is slower and when that is the case the performance disadvantage is marginal. Per Table 1, JOB query 11b and SSB query Q3.4, form the worst results for TTJ are just 10% and 20% slower respectively. The weighted average of TTJ execution time over the three benchmarks is a hair better than 10% faster. More sizable improvements appear in the maximum speed-up results. We remind the reader the join orders are for plans that were optimized for left-deep linear hash-joins. Below we will return to the question of the upside opportunity for TTJ execution speed when, in future work, a SQL optimizer includes cost models for TTJ and the optimization process includes both a choice of join order and a choice of join algorithm. For instance the detailed examination of each query execution revealed that TTJ's worst-relative performance, JOB query 11b is due to the inclusion of the no-good list optimization which, often predictably, incures overhead without providing any performance benefit.

For completeness, performance of TTJ relative to YA is presented in Figure 6c. The results exemplify the paradox and challenge of YA. On all but 12 queries, TTJ outperforms YA, with average and maximum speedup of 1.4x and 9.2x. 8 of those queries, JOB 6a, 6b, 6c, 6d, 6e, 7b, 12b, and TPC-H Q7, exhibit the most significant disadvantage of TTJ.

Review of the JOB queries reveals a foreseeable cause for YA execution speed advantage. The first semijoin removes a large fraction of tuples from a large relation. For example, the first semijoin for JOB query 6a reduces the largest relation, `cast_info`, from 36,000,000 tuples to 486 tuples. That semijoin is executed before building the hash tables. Hash table build time for YA is 499ms. For TTJ that build time is 13,398ms and by itself comprises 98% of the execution time for TTJ.

The basis of TPC-H, Q7's performance results are also due to the impact of the first semijoin, but in a more involved way. Prior to any join processing a relational select on nation returns just 1 tuple. As an argument to the first semijoin, `supplier`  $\bowtie$  `nation`, over 90% of tuples from `supplier` are removed. Where, in the first example the one semijoin reduction accounted for speed benefit, in this example, by beginning with a single tuple, the entire chain of semijoin reductions resulted in large reductions in the size of the join arguments repeat.

Review of hash table build times for YA relative to hash table build times for TTJ and HJ alone, (these latter two always being equal), accounts for all the speed improvement of YA compared to the other algorithms.

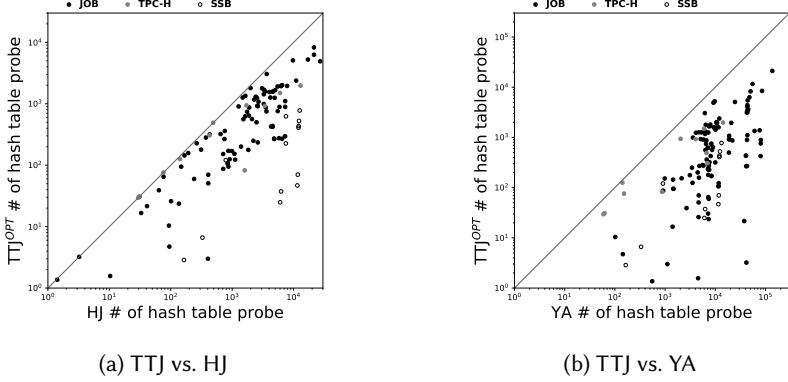


Fig. 7. Number of hash probes in different algorithms.

The specialized pattern embodied in SSB, star queries on a star schema, enables a quantitative assessment that may be used in the future by a query optimizer. Notably a determination if the integration of a no-good list is advantageous. For the special case of big data queries modeled by SSB, the performance of TTJ is largely determined by the effectiveness of the no-good list.

Recall the no-good list is a specialization for the leftmost argument of a join plan as hash-joins do not typically create a hash-table for the leftmost argument. The no-good list forms a cache of the tuple key values for the leftmost argument that have been determined to be dangling. Queries plans for star schema typically start with the fact table as the leftmost relation in a plan, and provide the key values for a series of joins on dimension tables. Any lookup failure will backjump to the fact table and add to the no-good list. Thus the no-good list acts as a filter that prevents any processing of a fact tuple whose join key values have already been determined to be fruitless.

Of 13 SSB queries evaluated TTJ is the fastest algorithm on 6, and, plus or minus, within 10% margin of the best algorithm on 10 queries. We compared the queries that run relatively slower in TTJ (Q1.2, Q3.4, Q4.1, and Q4.3) with those that run relatively faster (Q2.1, Q2.2, Q2.3, Q3.1, Q3.2, and Q3.3), and measured the ratio between the intermediate result size reduction and the size of no-good lists. We determined that for the slower queries, each element in the no-good list, on average, reduces the intermediate result size by 182. For the faster queries the average is 318. Although an optimizer is not within the scope of this paper, we can conclude that for our testbed more refined measurements would determine a tipping point value of a selectivity that falls between 1/182 and 1/318. Selectivity below the tipping point indicates omitting the no-good list will result in faster query execution and vice versa.

The scatter plot Figure 7a compares the number of hash probes for TTJ vs. HJ for each query. A small number of the scatter plot points appear on the diagonal, i.e. an equal number of hash probes. The remainder of the points are below the diagonal. This empirically validates our theorem that TTJ will execute fewer or an equal number of hash probes as HJ.

We have not made any claims as to the relative number of hash probes between TTJ and YA. Nevertheless we made that measurement. Figure 7b shows like HJ, TTJ makes fewer hash probes for YA. Yet YA runs faster for certain queries as hash building, not probing, sometimes dominates query run time, as we have pointed out in the analysis of results in Figure 6c.

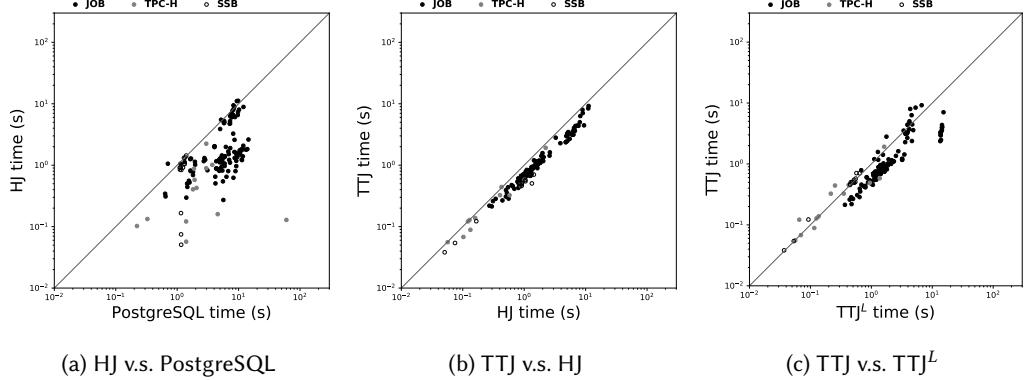


Fig. 8. Run time comparison of TTJ,  $TTJ^L$ , HJ and PostgreSQL.  $TTJ^L$  uses left-deep linear plans generated by SQLite (same as earlier experiments), while all other algorithms use bushy plans generated by PostgreSQL.

## 6.2 TTJ on Bushy Plans

Since every plan in Section 6.1 is compatible with a GYO reduction order, the runtime of TTJ is guaranteed to be linear. However, given the abundance of bushy plans (all the native PostgreSQL plans we used here are bushy), a natural question to ask is whether TTJ can still provide reasonably good performance despite the loss of linear runtime guarantee. The results in this section give an affirmative answer.

Figures 8a and 8b compare the run time of TTJ, HJ, and PostgreSQL using bushy plans produced by PostgreSQL, and Figure 8c compares the run time of TTJ using bushy plans with the same algorithm using left-deep linear plans produced by SQLite, denoted by  $TTJ^L$ . Figure 8a shows our HJ baseline remains competitive with PostgreSQL under bushy plans. TTJ is faster than HJ on all 113 JOB queries, and faster than  $TTJ^L$  on 101 (89%) of them. Compared to  $TTJ^L$ , the maximum speed-up is  $5.7\times$  (7b), the minimum speed-up is  $0.5\times$  (19d), and the average speed-up (geometric mean) is  $1.75\times$ . Compared to HJ, the maximum speed-up is  $2.1\times$  (13d), the minimum speed-up is  $1.1\times$  (19d), and the average speed-up (geometric mean) is  $1.56\times$ . From the figure, we observe that the materialization of intermediate results due to bushy plans does not degrade TTJ performance; in fact, TTJ performs much better than itself on linear plans in most cases due to the fact that intermediate results generated in bushy plans are usually smaller than some of the largest input relations in JOB, which allows TTJ to spend less time building hash tables. As a result, the saving in join computation becomes more salient than that over left-deep plans. Furthermore, the performance improvement using TTJ on each individual linear plan has the compound effect that contribute to the overall performance improvement of the queries. However, we do observe that some of the queries still have better performance under linear plans than bushy plans such as 8c and 16b. This indicates that linear time guarantee is still meaningful for query performance and optimization is still necessary to bring out the best performance of TTJ (e.g., decide which plan shape to use). An important topic for future work is to optimize bushy plans that also have the guarantee of optimality.

## 6.3 Implementation-level Comparison

Although TTJ offers strong theoretical guarantees and performs well in the controlled experiments in Sections 6.1 and 6.2, it is not competitive with state-of-the-art systems. To achieve stronger performance, we have developed the  $TTJ^\epsilon$  variant of our algorithm, as described in Section 5. In this section, we compare  $TTJ^\epsilon$  with DuckDB [25], a state-of-the-art in-memory analytical database, as

Baseline	Average	Maximum	Minimum
DuckDB	3.83×	35.26× (4c)	0.81× (33c)
RPT	2.67×	11.30× (3b)	0.70× (17e)
YA <sup>+</sup>	5.42×	30.22× (2b)	1.28× (27b)

Table 2. Speed-up of  $\text{TTJ}^\epsilon$  Relative to DuckDB, Robust Predicate Transfer (RPT), and Yannakakis<sup>+</sup> (YA<sup>+</sup>).

well as two recent practical instantiations of YA: Robust Predicate Transfer (RPT) by Zhao et al. [40] and Yannakakis<sup>+</sup> (YA<sup>+</sup>) by Wang et al. [34]. RPT is implemented on top of DuckDB, while YA<sup>+</sup> works by rewriting the input SQL query and is compatible with any underlying database system. For consistency, we use the same DuckDB version (v0.9.2) used by Zhao et al. [40] for all baseline systems. We focus on the Join Order Benchmark in this section. In their experiments, Wang et al. [34] scaled up the input data by 10 to 100 times<sup>7</sup>. This introduced additional many-to-many joins, and resulted in more pronounced speedup by YA<sup>+</sup>. For consistency, we supply the original input data without change to all systems.

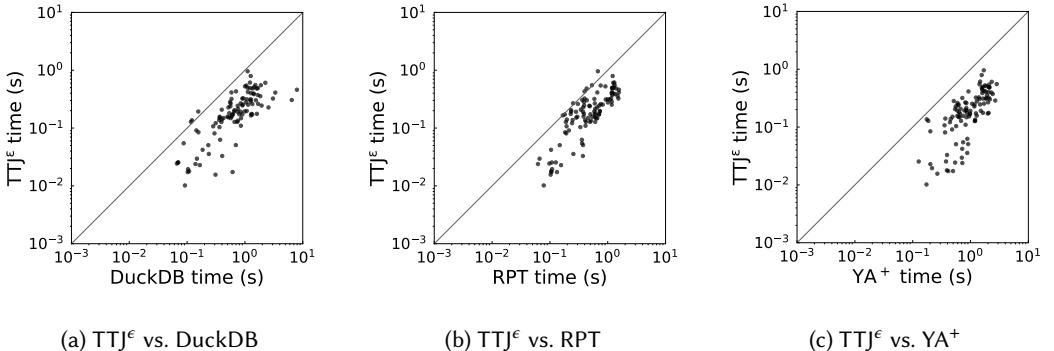


Fig. 9. Run time of  $\text{TTJ}^\epsilon$ , RPT, YA<sup>+</sup>, and DuckDB on JOB. Every data point corresponds to a query, whose  $x$ - and  $y$ -coordinates correspond to the run time of the algorithms under comparison. A point below the diagonal indicates that the algorithm on the  $y$ -axis is faster than the one on the  $x$ -axis.

Figure 9 and Table 2 show the experiment results. Appendix C contains additional comparisons of TTJ with the baselines. Overall,  $\text{TTJ}^\epsilon$  matches or outperforms all three baselines. The comparison with RPT and YA<sup>+</sup> show a slight trend of greater speedup for smaller (short-running) queries. This demonstrates the impact of TTJ<sup>ε</sup>'s zero-overhead guarantee: smaller queries finish quickly, so any overhead becomes more pronounced. But the impact of overhead also carries over to larger queries. Query 30b joins together 12 relations (with two self joins), and TTJ<sup>ε</sup> is 2.8× faster than RPT and 4.2× faster than YA<sup>+</sup>. This is because that query involves 6 large relations each containing several million tuples, totalling over 60 million rows. Although RPT's implementation of semijoin based on bloom

<sup>7</sup>Wang et al. [34] suggested that due to the prevalence of primary key/foreign key constraints in JOB, regular binary hash join can already achieve linear time performance, so there is little to gain from YA-style algorithms. However, we have observed superlinear behavior in practice: even if *all* joins are on primary/foreign keys, the same primary key may be joined with multiple foreign keys, and the optimizer may choose (perhaps based on bad cardinality estimate) to join foreign keys with each other first, leading to superlinear blowup.

filters is very fast, it still requires two semijoin passes in general<sup>8</sup>, incurring significant overhead. The overhead for  $YA^+$  comes from two sources. First,  $YA^+$  works by rewriting the input SQL query. To ensure the underlying database follows the prescribed query plan,  $YA^+$  creates one temporary view per operation. According to Wang et al. [34] this carries an overhead, and for query 30b the rewritten query creates 22 views. Second, when the final result is aggregated  $YA^+$  aggressively pushes down aggregation to eliminate the need for semijoin reduction. This guarantees linear time complexity, but may introduce additional overhead. For example, when each group-by value has low degree (i.e., they appear in a small number of tuples), early aggregation cannot effectively reduce the relation size but still adds a full scan. And if a large number of groups are eventually removed by a later join, the aggregation work on them would be wasted. The rewritten plan by  $YA^+$  for query 30b contains 25  $\text{MIN}$  aggregates, while the original query has only 4.

There are a few queries where  $TTJ^e$  is slightly slower than the baseline compared. The reason is mostly the inaccuracy of the cardinality estimation of our optimizer for  $TTJ^e$ , which leads to suboptimal plans. For example, in 27b,  $TTJ^e$  has a speed-up 0.89 $\times$  over DuckDB. In this query, `movie_keyword` is the largest with around 4.5 million tuples. After joining with `keyword`, the relation has around 10 thousands remaining tuples, which is the smallest among all the relations joining on `movie_id`. DuckDB optimizer estimates that fact accurately and evaluates this join first, which effectively uses it as a filter to prune out dangling tuples from other relations and saves the hash-table build time during the query evaluation. On the other hand, our optimizer estimates the join result to be the second smallest, which misses some opportunities to prune dangling tuples early. Contrast these observations with the findings by Zhao et al. [40]: they found RPT to have similar performance regardless of the query plan, concluding the algorithm to be robust. This is because the semijoin in RPT is based on bloom filters, so each probe has a constant cost. And since RPT performs two full semijoin passes, which involves scanning each relation and probing into the neighboring bloom filters, there is no significant variation in run time. During the final join phase, there is also a fixed cost per output tuple, so the overall run time is predictable. In contrast,  $TTJ^e$  eliminates the second semijoin pass and directly probes into the hash tables without relying on additional filter structures. On one hand, this leads to lower overhead, and on the other  $TTJ^e$  is more sensitive to the query plan. If stability is desirable, bloom filters can also be added to  $TTJ^e$  to trade off strong best-case performance to improve the average case.

## 7 Cyclic Queries

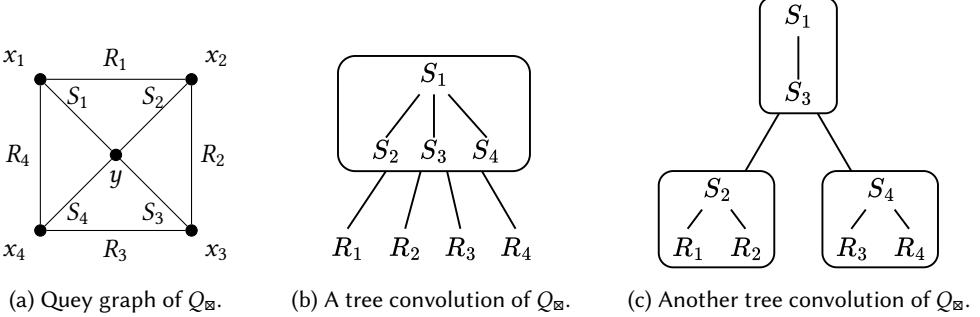
The  $TTJ$  algorithm as defined in Figure 4a supports both acyclic and cyclic queries. The guarantee to match or outperform binary join (Theorem 4.10) also holds for cyclic queries. However, the linear-time guarantee only applies to acyclic queries. To analyze the run time of  $TTJ$  on cyclic queries, we introduce a new method called *tree convolution* to break down a cyclic query in to acyclic parts. The next example illustrates the intuition behind tree convolution.

*Example 7.1.* Consider the following query whose query graph is shown in Figure 10a:

$$\begin{aligned} Q_{\bowtie} : -R_1(x_1, x_2) \bowtie R_2(x_2, x_3) \bowtie R_3(x_3, x_4) \bowtie R_4(x_4, x_1) \bowtie \\ S_1(x_1, y) \bowtie S_2(x_2, y) \bowtie S_3(x_3, y) \bowtie S_4(x_4, y) \end{aligned}$$

In the query graph, each node is a variable, and there is an edge between two nodes if the corresponding variables appear in the same atom, for example the edge  $R_1$  between  $x_1$  and  $x_2$  corresponds to the atom  $R_1(x_1, x_2)$ . Clearly the query is cyclic. However, we can compute the query by breaking it down into two acyclic steps: First, we compute the join  $S_1 \bowtie S_2 \bowtie S_3 \bowtie S_4$  and store the result in a temporary relation  $S$ . Then we compute the final result with  $R_1 \bowtie R_2 \bowtie R_3 \bowtie R_4 \bowtie S$ . Any acyclic

<sup>8</sup>Only when the join phase plan coincides with the second semijoin phase plan, RPT skips the second semijoin phase.

Fig. 10. Query graph and tree convolutions of  $Q_{\otimes}$  in Example 7.1.

join algorithm can be used to compute each step, and, as we will show later, using TTJ can avoid materializing the intermediate result  $S$ . The total run time is therefore  $O(|\text{IN}| + |\text{OUT}| + |S|)$ .

To formally define tree convolutions, we first identify a (full conjunctive) query with the set of atoms in its body, and define a *subquery* of a query  $Q$  as a subset of the atoms in  $Q$ . Then, a tree convolution for a query is a nested tree, defined recursively as follows:

*Definition 7.2.* A tree convolution of a query  $Q$ , written  $\mathcal{T}(Q)$ , is a tree where each node is either an atom in  $Q$ , or a tree convolution of a subquery of  $Q$ . Each atom in  $Q$  appears exactly once anywhere in  $\mathcal{T}(Q)$ , and every tree  $T$  in  $\mathcal{T}(Q)$  forms a join tree of the corresponding (sub-)query, after replacing each non-atom node  $v \in T$  with a fresh atom over all variables in  $v$ .

Note that the recursive definition above will lead to nested trees, because each tree node can consist of a tree convolution of a subquery. Figure 10b shows a tree convolution corresponding to the computation of  $Q_{\otimes}$  in Example 7.1: first join together the  $S$  relations, then join the result with the  $R$  relations. Figure 10c shows a different convolution, where we first compute three acyclic subqueries, then join together the results. Another convolution with four nesting levels is shown in Figure 11b.

We can compute any cyclic query using a tree convolution. Starting from the most deeply nested trees, we run any acyclic join algorithm to materialize intermediate results, until we reach the top-level tree that produces the final output. However, TTJ can avoid the expensive materialization by using a special kind of tree convolution called *rooted convolution*.

*Definition 7.3.* A convolution is *rooted* if nested convolutions only appear at the root of each tree.

The convolution in Figure 10b is rooted, while the one in Figure 10c is not. We can generate a plan  $p$  for TTJ by traversing a rooted convolution inside-out: starting with the most deeply nested tree, initialize  $p$  with the reverse of the GYO-reduction order of this tree; then, as we go up each level, append the reverse of the GYO-reduction order to the end of  $p$ .

*Example 7.4.* The rooted convolution in Figure 10b generates the plan  $[S_1, S_2, S_3, S_4, R_1, R_2, R_3, R_4]$ .

A small adjustment to the TTJ algorithm is necessary to fully exploit rooted convolutions. If we execute TTJ as-is using the plan in Example 7.4, none of the relations  $R_1, \dots, R_4$  have a parent in the plan, yet we need to compute the join  $S \bowtie R_1 \bowtie \dots \bowtie R_4$  in time  $O(|S| + \sum_i |R_i| + |Q|)$ , where  $S$  is the join of  $S_1, \dots, S_4$ . We therefore introduce additional backjumps from each  $R_i$  to  $S_4$ , but without deleting any tuple from  $S_4$ . This is achieved by defining the parent function to work over rooted convolutions: given a tree convolution  $C$  and a relation  $R$ , if the parent node of  $R$  in some tree of  $C$

is an atom  $P$ , then assign  $P$  as the parent of  $R$ ; otherwise if the parent node is a nested tree, then assign the last relation in that tree (i.e. the first relation in its GYO-reduction order) as the parent of  $R$ . For example, the parent of each  $R_i$  in Example 7.4 is  $S_4$ .

Finally, when backjumping to a parent in a nested convolution, we do not delete any tuple from that parent. We are now ready to analyze the run time of TTJ on cyclic queries.

Given a rooted tree convolution  $\mathcal{T}(Q)$ , we generate a query plan  $p$  as follows. Suppose  $\mathcal{T}(Q)$  consists of  $m$  nested trees  $T_1, \dots, T_m$  where  $T_1$  is the most deeply nested tree and  $T_m$  is the outermost tree. Let  $p_1, \dots, p_m$  be the plans corresponding to  $T_1, \dots, T_m$ . Then, each  $p_i$  is a plan that corresponds to the reverse of a GYO-reduction order of  $T_i$ , where the first relation in  $p_i$  is result of  $p_{i-1}$  for  $i > 1$ . The outermost plan  $p_m$  then computes the final result of  $Q$ .

**PROPOSITION 7.5.** *During ttj execution on a given rooted convolution  $\mathcal{T}$  of a query  $Q$ , if a lookup fails at  $R$  that belongs to  $p_i$  but not in  $p_{i-1}$  ( $i \geq 2$ ), ttj either backjumps to an atom that is in  $p_i$  but not in  $p_{i-1}$  or backjumps to the last atom of  $p_{i-1}$ .*

**PROOF.** For any  $p_i$  with  $i \geq 2$ , since  $p_i$  is the reverse of a GYO-reduction order of the  $i$ -th tree of  $\mathcal{T}$  and by the modified parent function, every relation in  $p_i$  but not in  $p_{i-1}$  has parent. Furthermore, if lookup fails at  $R$  that is in  $p_i$  but not in  $p_{i-1}$ , ttj backjumps to  $R$ 's parent. Then, the result follows by the definition of the modified parent function.  $\square$

In  $p_1$ , if the parent of a relation is the first relation of  $p_1$ , ttj backjumps to the first relation. We can treat the first relation of  $p_1$  as  $p_0$  (i.e., the most deeply nested tree in  $\mathcal{T}$  is now a node of an atom of  $Q$ ) and the relation is also the last atom of  $p_0$ . Therefore, we can remove the restriction of  $i \geq 2$  in Proposition 7.5. In the following proof, we reference the Proposition 7.5 with the understanding that it holds for  $i \geq 1$ .

**THEOREM 7.6.** *Given a rooted convolution  $\mathcal{T}(Q)$  of  $Q$ , there is a plan  $p$  such that TTJ runs in time  $O(|IN| + |OUT| + \sum_i |S_i|)$  on  $p$  where  $|S_i|$  is the size of the join of all relations in the  $i$ -th tree of  $\mathcal{T}(Q)$ .*

**PROOF.** Since there is no change to ttj except using the modified parent function for  $\mathcal{T}$ , like Proof of Theorem 4.7, we only need to show the algorithm makes  $O(|IN| + |OUT| + \sum_i |S_i|)$  number of calls to ttj.

Except for the initial call to ttj with  $ttj(\emptyset, p, 1)$ , every call to ttj has three possible outcomes: (1) It outputs a tuple. (2) It backjumps and possibly deletes a tuple from an input relation. (3) It recursively calls ttj. Because the query plan has constant length, there can be at most a constant number of recursive calls to ttj (case 3) until we reach cases 1 or 2. There can be  $O(|Q|)$  ttj calls for case 1. Since lookup cannot fail at the first relation of  $p$ , the relation that lookup fails at is in some  $p_i$  but not in  $p_{i-1}$ . Let  $R$  be a relation that a lookup fails at. By Proposition 7.5, there can be two cases on where ttj backjumps to. If ttj backjumps to an atom that is also in  $p_i$  but not in  $p_{i-1}$ , a tuple is deleted. This case can happen  $O(|IN|)$  times. If ttj backjumps to the last atom of  $p_{i-1}$ , since ttj works no different than binary join from this moment until next lookup failure, this can happen  $O(|S_{i-1}|)$  times; the tuples computed at the last atom of  $p_{i-1}$  is  $S_{i-1}$ . Therefore, given  $1 \leq i \leq m$ , there are at most  $O(|IN| + |OUT| + \sum_{i=1}^{m-1} |S_i|)$  calls to ttj, and the algorithm runs in that time.  $\square$

We conclude this section by noting that tree convolution generalizes several existing ideas in databases and constraint satisfaction. First, classic binary join plans are a special case, where every tree in the tree convolution is of size 2. For example, the convolution in Figure 11b corresponds to the binary join plan in Figure 11a. In other words, traditional hash join can be thought of as computing a convolution one binary join at a time, and we have generalized this to computing a multi-way acyclic join at a time. In a similar way, rooted convolutions generalize left-deep linear

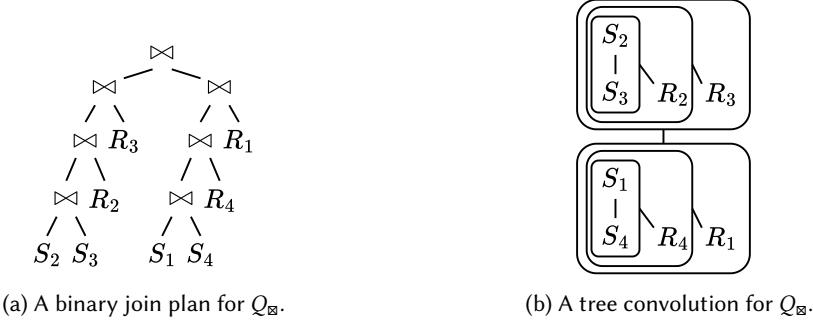


Fig. 11. A binary join plan for the query in Example 7.1 and the corresponding tree convolution.

plans, as the top-half of Figure 11b corresponds to the left-half in Figure 11a. Second, in constraint satisfaction a *cycle cut set* is a subset of the constraints whose exclusion makes the constraint problem acyclic. In database terms, it is a subset of the atoms of  $Q$  whose removal leaves an acyclic subquery. Tree convolution generalizes cycle cut sets in the sense that it “cuts” the query in multiple rounds, with each round producing acyclic subqueries that can be computed in linear time. Finally, many of the properties of tree convolution are shared with tree decomposition of hypergraphs [10] in database theory. One difference is that tree convolution allows nesting, while tree decomposition is a flat partitioning of the query graph. Another difference is that tree decomposition usually requires different algorithms for joining within each bag (e.g. worst-case optimal join [23]) and for joining the bags together (e.g. YA), while tree convolution uses the same algorithm for all nesting levels. This is enabled precisely by the nesting structure: at every nesting level the subquery forms a tree, and can be computed by an acyclic join algorithm. Finally, while complexity bounds derived from tree decompositions are known to be tight [11], further research is needed to study the tightness of tree convolution-based bounds.

## 8 Future Work and Conclusion

In this paper we have proposed our new join algorithm, TreeTracker Join (TTJ). The algorithm runs in time  $O(|\text{IN}| + |\text{OUT}|)$  on acyclic queries, and guarantees to make no more hash probes than binary hash join on the same query plan. We have shown empirically that TTJ is competitive with binary hash join and Yannakakis’s algorithm.

Although our implementation already beats PostgreSQL in our experiments, challenges remain for TTJ to compete with highly optimized systems. Decades of research on binary join has produced effective techniques like column-oriented storage, vectorized execution, and parallel execution, just to name a few. In the following, we outline several directions for future work.

*Query Optimization.* Another avenue for future work is to develop a dedicated query optimizer for TTJ. As this paper’s focus is on algorithm-level comparison, we have opted to reuse existing systems to produce binary hash join plans, which are then executed using TTJ. Tailoring the optimizer to TTJ may yield plans with better performance. For instance, estimating the number of hash probe failures instead of intermediate result sizes shall more accurately model the execution cost of TTJ. On the other hand, extending TTJ to also guarantee a linear time complexity on bushy plans is also an interesting challenge. Our theoretical analyses of TTJ reveal a close connection between GYO-reduction orders and left-deep linear plans, both of which are total orders. Since bushy plans and join trees both define partial orders, we conjecture there exists a algorithm that

runs in linear time on any bushy plan that corresponds to a join tree, with the same guarantee of matching binary hash join on the same plan.

*Multi-threading.* Nearly all modern database systems are multi-threaded. While we have focused on single-threaded execution in this paper, making TTJ multi-threaded brings unique challenges. The current design of TTJ interleaves deletion and probes on hash tables, and a naïve parallelization using locks would lead to bad performance. An alternative is to tag each tuple with a boolean *tombstone* that is atomically set whenever that tuple is deleted. The TTJ<sup>c</sup> variant of the algorithm also leverages query structure and primary/foreign key constraints to avoid unnecessary deletions, which should make the algorithm more amenable to parallelization.

*Cyclic Queries.* Our experiments focus on acyclic queries due to their prevalence in traditional workloads. However, with the rise of graph databases practitioners begin to encounter more and more cyclic queries. Additional research on TTJ for cyclic queries, both in terms of practical performance and theoretical guarantees, will be very valuable. Some open problems include: Given any hypergraph, what is the minimum nesting depth of any tree convolution? How does this number relate to other measures like various notions of hypergraph widths? And what is the complexity of finding the optimal tree convolution given a cost function? Answering these questions will aid the development of a query optimizer for TTJ on cyclic queries.

*Integration with Existing Systems.* Bringing ideas from the lab into production requires significant engineering effort. While a full integration of TTJ into a mainstream database is out of the scope of this paper, we discuss challenges in doing so and sketch a blueprint for an implementation in SQLite in Appendix B.

## References

- [1] [n. d.]. Java Microbenchmark Harness (JMH). <https://github.com/openjdk/jmh>
- [2] Guillaume Bagan, Arnaud Durand, and Etienne Grandjean. 2007. On Acyclic Conjunctive Queries and Constant Delay Enumeration. In *Annual Conference for Computer Science Logic*. <https://api.semanticscholar.org/CorpusID:15398587>
- [3] Guillaume Bagan, Arnaud Durand, and Etienne Grandjean. 2007. On Acyclic Conjunctive Queries and Constant Delay Enumeration. In *Computer Science Logic, 21st International Workshop, CSL 2007, 16th Annual Conference of the EACSL, Lausanne, Switzerland, September 11-15, 2007, Proceedings (Lecture Notes in Computer Science, Vol. 4646)*, Jacques Duparc and Thomas A. Henzinger (Eds.). Springer, 208–222. doi:10.1007/978-3-540-74915-8\_18
- [4] Liese Bekkers, Frank Neven, Stijn Vansumeren, and Yisu Remy Wang. 2024. Instance-Optimal Acyclic Join Processing Without Regret: Engineering the Yannakakis Algorithm in Column Stores. *CoRR* abs/2411.04042 (2024). doi:10.48550/ARXIV.2411.04042 arXiv:2411.04042
- [5] Altan Birler, Alfons Kemper, and Thomas Neumann. 2024. Robust Join Processing with Diamond Hardened Joins. *Proc. VLDB Endow.* 17, 11 (2024), 3215–3228. doi:10.14778/3681954.3681995
- [6] Dean Daniels. 1982. *Query compilation in a distributed database system*. Ph. D. Dissertation. Massachusetts Institute of Technology, Department of Electrical Engineering ....
- [7] Rina Dechter. 1990. Enhancement Schemes for Constraint Processing: Backjumping, Learning, and Cutset Decomposition. *Artif. Intell.* 41, 3 (1990), 273–312. doi:10.1016/0004-3702(90)90046-3
- [8] Kevin P. Gaffney, Martin Prammer, Laurence C. Brasfield, D. Richard Hipp, Dan R. Kennedy, and Jignesh M. Patel. 2022. SQLite: Past, Present, and Future. *Proc. VLDB Endow.* 15, 12 (2022), 3535–3547. doi:10.14778/3554821.3554842
- [9] César A. Galindo-Legaria, Torsten Grabs, Sreenivas Gukal, Steve Herbert, Aleksandras Surna, Shirley Wang, Wei Yu, Peter Zabback, and Shin Zhang. 2008. Optimizing Star Join Queries for Data Warehousing in Microsoft SQL Server. In *Proceedings of the 24th International Conference on Data Engineering, ICDE 2008, April 7-12, 2008, Cancún, Mexico*, Gustavo Alonso, José A. Blakeley, and Arbee L. P. Chen (Eds.). IEEE Computer Society, 1190–1199. doi:10.1109/ICDE.2008.4497528
- [10] Georg Gottlob, Gianluigi Greco, Nicola Leone, and Francesco Scarello. 2016. Hypertree Decompositions: Questions and Answers. In *Proceedings of the 35th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems (San Francisco, California, USA) (PODS '16)*. Association for Computing Machinery, New York, NY, USA, 57–74. doi:10.1145/2902251.2902309
- [11] Georg Gottlob, Gianluigi Greco, Nicola Leone, and Francesco Scarello. 2016. Hypertree Decompositions: Questions and Answers. In *Proceedings of the 35th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*,

*PODS 2016, San Francisco, CA, USA, June 26 - July 01, 2016*, Tova Milo and Wang-Chiew Tan (Eds.). ACM, 57–74. doi:10.1145/2902251.2902309

[12] M. Graham. 1980. *On the universal relation*. Technical Report. University of Toronto, Computer Systems Research Group.

[13] Muhammad Idris, Martín Ugarte, and Stijn Vansumeren. 2017. The Dynamic Yannakakis Algorithm: Compact and Efficient Query Processing Under Updates. In *Proceedings of the 2017 ACM International Conference on Management of Data, SIGMOD Conference 2017, Chicago, IL, USA, May 14–19, 2017*, Semih Salihoglu, Wenchao Zhou, Rada Chirkova, Jun Yang, and Dan Suciu (Eds.). ACM, 1259–1274. doi:10.1145/3035918.3064027

[14] Zachary G. Ives and Nicholas E. Taylor. 2008. Sideways Information Passing for Push-Style Query Processing. In *Proceedings of the 24th International Conference on Data Engineering, ICDE 2008, April 7–12, 2008, Cancún, Mexico*, Gustavo Alonso, José A. Blakeley, and Arbee L. P. Chen (Eds.). IEEE Computer Society, 774–783. doi:10.1109/ICDE.2008.4497486

[15] Roberto J. Bayardo Jr. and Daniel P. Miranker. 1994. An Optimal Backtrack Algorithm for Tree-Structured Constraint Satisfaction problems. *Artif. Intell.* 71, 1 (1994), 159–181. doi:10.1016/0004-3702(94)90064-7

[16] Phokion G. Kolaitis and Moshe Y. Vardi. 2000. Conjunctive-Query Containment and Constraint Satisfaction. *J. Comput. Syst. Sci.* 61, 2 (2000), 302–332. doi:10.1006/JCSS.2000.1713

[17] Kukjin Lee, Anshuman Dutt, Vivek R. Narasayya, and Surajit Chaudhuri. 2023. Analyzing the Impact of Cardinality Estimation on Execution Plans in Microsoft SQL Server. *Proc. VLDB Endow.* 16, 11 (2023), 2871–2883. doi:10.14778/3611479.3611494

[18] Viktor Leis, Andrey Gubichev, Atanas Mirchev, Peter Boncz, Alfons Kemper, and Thomas Neumann. 2015. How Good Are Query Optimizers, Really? *Proc. VLDB Endow.* 9, 3 (Nov. 2015), 204–215. doi:10.14778/2850583.2850594

[19] Lothar F. Mackert and Guy M. Lohman. 1986. R\* Optimizer Validation and Performance Evaluation for Distributed Queries. In *VLDB'86 Twelfth International Conference on Very Large Data Bases, August 25–28, 1986, Kyoto, Japan, Proceedings*, Wesley W. Chu, Georges Gardarin, Setsuo Ohsuga, and Yahiko Kambayashi (Eds.). Morgan Kaufmann, 149–159. <http://www.vldb.org/conf/1986/P149.PDF>

[20] Daniel P. Miranker, Roberto J. Bayardo, and Vasilis Samoladas. 1997. Query Evaluation as Constraint Search; An Overview of Early Results. In *International Symposium on the Applications of Constraint Databases*. <https://api.semanticscholar.org/CorpusID:8644835>

[21] Thomas Neumann. 2011. Efficiently Compiling Efficient Query Plans for Modern Hardware. *Proc. VLDB Endow.* 4, 9 (2011), 539–550. doi:10.14778/2002938.2002940

[22] Hung Q. Ngo, Ely Porat, Christopher Ré, and Atri Rudra. 2018. Worst-Case Optimal Join Algorithms. *J. ACM* 65, 3, Article 16 (March 2018), 40 pages. doi:10.1145/3180143

[23] Hung Q. Ngo, Ely Porat, Christopher Ré, and Atri Rudra. 2018. Worst-case Optimal Join Algorithms. *J. ACM* 65, 3 (2018), 16:1–16:40. doi:10.1145/3180143

[24] Pat O’Neil, Betty O’Neil, and Xuedong Chen. 2009. *Star Schema Bechmark - Revision 3, June 5, 2009*. resreport. UMass/Boston. <https://www.cs.umb.edu/~poneil/StarSchemaB.PDF>

[25] Mark Raasveldt and Hannes Mühlleisen. 2019. DuckDB: An Embeddable Analytical Database. In *Proceedings of the 2019 International Conference on Management of Data* (Amsterdam, Netherlands) (SIGMOD ’19). Association for Computing Machinery, New York, NY, USA, 1981–1984. doi:10.1145/3299869.3320212

[26] Praveen Seshadri, Joseph M. Hellerstein, Hamid Pirahesh, T. Y. Cliff Leung, Raghu Ramakrishnan, Divesh Srivastava, Peter J. Stuckey, and S. Sudarshan. 1996. Cost-Based Optimization for Magic: Algebra and Implementation. In *Proceedings of the 1996 ACM SIGMOD International Conference on Management of Data, Montreal, Quebec, Canada, June 4–6, 1996*, H. V. Jagadish and Inderpal Singh Mumick (Eds.). ACM Press, 435–446. doi:10.1145/233269.233360

[27] SQLite Documentation. 2024. Query Planning and Optimization. [https://www.sqlite.org/optoverview.html#hash\\_joins](https://www.sqlite.org/optoverview.html#hash_joins). Accessed: 2024-07-24.

[28] SQLite Documentation Team. 2025. The SQLite Bytecode Engine. <https://www.sqlite.org/opcode.html>. Accessed: 2025-08-21.

[29] Robert Endre Tarjan and Mihalis Yannakakis. 1984. Simple Linear-Time Algorithms to Test Chordality of Graphs, Test Acyclicity of Hypergraphs, and Selectively Reduce Acyclic Hypergraphs. *SIAM J. Comput.* 13, 3 (1984), 566–579. doi:10.1137/0213035

[30] Transaction Processing Performance Council (TPC). [n. d.]. TPC-H Benchmark. Online. [http://tpc.org/tpc\\_documents\\_current\\_versions/pdf/tpc-h\\_v3.0.pdf](http://tpc.org/tpc_documents_current_versions/pdf/tpc-h_v3.0.pdf) Accessed on 11-18-2021.

[31] Nikolaos Tziavelis, Wolfgang Gatterbauer, and Mirek Riedewald. 2022. Any-k Algorithms for Enumerating Ranked Answers to Conjunctive Queries. *CoRR* abs/2205.05649 (2022). doi:10.48550/ARXIV.2205.05649 arXiv:2205.05649

[32] Nikolaos Tziavelis, Wolfgang Gatterbauer, and Mirek Riedewald. 2024. Ranked Enumeration for Database Queries. *SIGMOD Rec.* 53, 3 (2024), 6–19. doi:10.1145/3703922.3703924

[33] Todd L. Veldhuizen. 2012. Leapfrog Triejoin: a worst-case optimal join algorithm. *CoRR* abs/1210.0481 (2012). arXiv:1210.0481 <http://arxiv.org/abs/1210.0481>

- [34] Qichen Wang, Bingnan Chen, Binyang Dai, Ke Yi, Feifei Li, and Liang Lin. 2025. Yannakakis+: Practical Acyclic Query Evaluation with Theoretical Guarantees. *Proc. ACM Manag. Data* 3, 3, Article 235 (June 2025), 28 pages. doi:10.1145/3725423
- [35] Qichen Wang, Xiao Hu, Binyang Dai, and Ke Yi. 2023. Change Propagation Without Joins. *Proc. VLDB Endow.* 16, 5 (2023), 1046–1058. doi:10.14778/3579075.3579080
- [36] Qichen Wang and Ke Yi. 2022. Conjunctive Queries with Comparisons. In *SIGMOD '22: International Conference on Management of Data, Philadelphia, PA, USA, June 12 - 17, 2022*, Zachary G. Ives, Angela Bonifati, and Amr El Abbadi (Eds.). ACM, 108–121. doi:10.1145/3514221.3517830
- [37] Yisu Remy Wang, Max Willsey, and Dan Suciu. 2023. Free Join: Unifying Worst-Case Optimal and Traditional Joins. *Proc. ACM Manag. Data* 1, 2 (2023), 150:1–150:23. doi:10.1145/3589295
- [38] Mihalis Yannakakis. 1981. Algorithms for Acyclic Database Schemes. In *Very Large Data Bases, 7th International Conference, September 9-11, 1981, Cannes, France, Proceedings*. IEEE Computer Society, 82–94.
- [39] Clement T. Yu and M. Z. Ozsoyoglu. 1979. An algorithm for tree-query membership of a distributed query. In *Annual International Computer Software and Applications Conference*. <https://api.semanticscholar.org/CorpusID:7812638>
- [40] Junyi Zhao, Kai Su, Yifei Yang, Xiangyao Yu, Paraschos Koutris, and Huanchen Zhang. 2025. Debunking the Myth of Join Ordering: Toward Robust SQL Analytics. *Proceedings of the ACM on Management of Data* 3, 3, Article 146 (June 2025), 28 pages. doi:10.1145/3725283
- [41] Jianqiao Zhu, Navneet Potti, Saket Saurabh, and Jignesh M. Patel. 2017. Looking Ahead Makes Query Plans Robust. *Proc. VLDB Endow.* 10, 8 (2017), 889–900. doi:10.14778/3090163.3090167

## A No-Good List

We had remarked in Section 4 that removing a tuple from the root relation is pointless, as the same tuple would never be considered again. However, any tuple in the root relation that shares the same values with an offending tuple over the key schema will also fail. The no-good list optimization comprises adding that set of values to a blacklist. Each tuple from the root relation is tested for membership in the blacklist. Since membership in that list mean certain failure no further effort to join that tuple is necessary. This optimization requires three changes to Figure 4a.

First, the key values must be included as parameters and passed to the parent relation, line 7:

```
return (P, πR(t))
```

When catching the backjump (line 10) at the root relation, those key values are added to the blacklist:

```
if result == (R, vals):
    if i == 0: no_good.add(vals) else: R[k].delete(r)
```

When iterating over the root relation, (after line 8), each tuple is tested for membership in the no-good list and if present further processing is skipped.

```
if i == 0 & r.matches(no_good): continue
```

*Example A.1.* Consider the join  $T(y, k) \bowtie U(y, l)$  with the relations as defined in Example 1.1. As we iterate over  $T$ , the first lookup  $U[y = 1]$  immediately fails. Upon this we add  $y = 1$  to the no-good list for  $T$ , then in subsequent iterations we need only check the no-good list instead of probing into  $U$ . When failures are concentrated on a few values, the no-good list will be small, and checking it can be faster than probing into a large hash table. Another case when the no-good list is useful is when the deletion happens after backjumping multiple levels. Then, checking the no-good list requires one single probe as opposed to probing into each level that was backjumped over.

The no-good list,  $ng$ , can be implemented as a hash table. Suppose the root relation,  $R$ , has  $m$  children  $S_1, \dots, S_m$ . The lookup key for  $ng$  is  $\langle S_i, \ell_i \rangle$  where  $\ell_i$  is a set containing  $\pi_{\text{keys}(R, S_i)}(t)$  (called *no-goods*) for a tuple  $t$  from  $R$  that caused a lookup failure at  $S_i$  for  $i \in [m]$ . The impact of the no-good list is almost identical to semijoin reduction in YA. The algorithmic difference is in lieu of a semijoin removing dangling tuples prior to the join, the  $R$  tuples are checked against a collection of values accumulated on the fly and at anytime during execution are a subset of the contents of the complementary antijoin. Like YA itself, the effectiveness of the no-good list depends on how much the argument is reduced and the size of the intermediate result. i.e. the semijoin and join selectivity.

## B SQLite Implementation

Two main challenges of implementing TTJ are the interruption of control flow by backjumping and the need to delete tuples while probing hash tables. After reviewing the source code of SQLite, we find it to be uniquely suited for integrating TTJ. First, SQLite's query execution engine is based on a bytecode virtual machine [28], where loops are already implemented with GOTO-like jumps. Backjumping in TTJ can therefore be implemented by analyzing the join tree structure to determine backjumping points, then generating the bytecode to perform the jumps upon hash lookup failures. Second, deletion during iteration may lead to undefined behavior without care. Fortunately, the documentation of SQLite's `DELETE` bytecode specifically calls out that “it is ok to delete a record from within a loop” [28]. The introduction of bloom filters in SQLite [8] also makes it possible to implement the TTJ<sup>e</sup> variant of our algorithm. Specifically, for each parent in a join tree we can simultaneously probe into the bloom filter of all child relations, and delete the current tuple from the

parent relation should any probe fail. One complication is that SQLite constructs temporary B-tree indices (standing in for hash tables) *lazily*, so early probing as in  $\text{TTJ}^\epsilon$  is not always possible. On one hand, to guarantee instance-optimality, early probing is crucial; on the other, lazy construction is only beneficial when the join result is empty, which is rarely the case for OLAP queries.

### C Performance Comparison of TTJ with $\text{TTJ}^\epsilon$ , RPT and $\text{YA}^+$

Figure 12 compares the performance of TTJ,  $\text{TTJ}^\epsilon$ , RPT and  $\text{YA}^+$  on JOB.

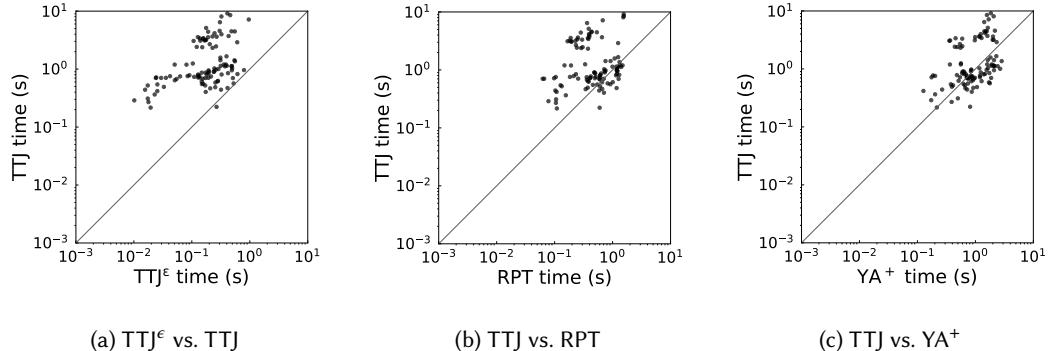


Fig. 12. Run time of TTJ,  $\text{TTJ}^\epsilon$ , RPT, and  $\text{YA}^+$  on JOB. Every data point corresponds to a query, whose  $x$ - and  $y$ -coordinates correspond to the run time of the algorithms under comparison. A point below the diagonal indicates that the algorithm on the  $y$ -axis is faster than the one on the  $x$ -axis.

We thank the referees for their thoughtful comments. In the following, we will first address some common concerns. Then, we address the comments of each referee in detail. Relevant changes to common concerns in the paper are highlighted in **red**.

### Comparison with recent work

Referee 2&3 requested additional experiments comparing with recent practical implementations of Yannakakis' algorithm. We added comparison with two most recent papers, Yannakakis+ (YA<sup>+</sup>) by Wang et al. [34] and Robust Predicate Transfer (RPT) by Zhao et al. [40]. Both papers appeared at SIGMOD 2025.

To make our algorithm competitive with these approaches, we developed a new variant of the TTJ algorithm which we call TTJ<sup>ε</sup>, described in Section 5. Despite its strong theoretical guarantees, the performance of the original TTJ algorithm is held back by two aspects. First, its reliance on backjumping complicates the control flow, resulting in complex code that is difficult for the compiler to optimize, and stands in the way of vectorization. Second, the interleaving of deletion and lookup also prevents optimizations applicable to read-only algorithms. The new TTJ<sup>ε</sup> variant completely eliminates backjumping and instead performs early probing. This is inspired by previous work on sideway information passing using bloom filters(bitmap filters. But unlike prior work, TTJ<sup>ε</sup> requires no additional filter data structure and retains the strong theoretical guarantees of TTJ, as we discuss in Section 5.1. Another inspiration is the work on Free Join by Wang et al. [37]. While Free Join is worst-case optimal (with the appropriate plan), TTJ and TTJ<sup>ε</sup> are instance-optimal for acyclic queries and guarantee to incur zero overhead compared to binary hash join. We also take inspiration from Yannakakis+ [34] and avoid deletion when possible by leveraging the query structure as well as primary/foreign key constraints, which we discuss in Section 5.2.

Section 6.3 contains the details of the comparison among the algorithms. At a high level, we see the three approaches (TTJ<sup>ε</sup>, Robust Predicate Transfer, Yannakakis+) as making different tradeoffs in the design space. On one end, Yannakakis+ is based on query rewriting and easily applicable to any underlying database system; on the other, TTJ<sup>ε</sup> requires implementing a new algorithm and achieves the best performance in our experiments. In the middle, Robust Predicate Transfer uses bloom filters, making it possible to modularly extend an existing system (DuckDB in their paper), while leading to performance between that of TTJ<sup>ε</sup> and Yannakakis+.

While the new TTJ<sup>ε</sup> variant shows strong performance, we decided to keep the comparison among the original TTJ, Hash Join, and Yannakakis' algorithm in Section 6.1 because it was done in a controlled setting where all algorithms are implemented from scratch in the same environment. A “textbook” implementation of hash join using the iterator model is used as the *control*, whereas YA and TTJ are implemented as simple extensions of the hash join implementation. The original experiments can be seen as algorithm-level comparisons and the new ones implementation-level comparisons

### Definition of tree convolution

Referee 2 and 3 expressed confusion around the definition of tree convolution. Perhaps the confusion arises from the fact that tree convolution is defined recursively. We added a clarification after Definition 7.2 to point out it will lead to nesting. Both referees also asked for the relationship between tree convolution and tree decomposition, and referee 2 correctly pointed out that a tree decomposition with only the root containing more than 1 relation can also be pipelined. However, the difference is that tree decomposition usually requires different algorithms for joining within the bag (e.g. worst-case optimal join) and joining the bags together (e.g. Yannakakis'), while tree convolution is only a conceptual device we use to analyze the behavior of TTJ, i.e., the same algorithm is used for joins nested at every level. This is enabled precisely by the nesting structure:

at every nesting level the subquery forms a tree, and can be computed by an acyclic join algorithm. Finally, while complexity bounds derived from tree decompositions are known to be tight, further research is needed to study the tightness of tree convolution-based bounds. We added these details to the end of Section 7.

## Response to Referee 1

Relevant changes in the paper are highlighted in **orange**.

*How can we safely eliminate tuples during multi-threaded execution?* We focus on the single-threaded execution in this paper, but we agree with the reviewer that multi-threaded execution introduces additional challenges. The new  $\text{TTJ}^\epsilon$  variant eliminates backjumping and avoids deletion when possible, which should make multi-threading easier. We added a paragraph in Section 8 to discuss several ideas for future research making  $\text{TTJ}$  multi-threaded.

*Handling cartesian products and non-equality joins.*  $\text{TTJ}$  is compatible with cartesian products and non-equality joins, and we added a discussion in Section 4.3. If the query contains a Cartesian product,  $\text{TTJ}$  would simply never backjump across the product operator, and the guarantees of instance-optimality and zero-overhead continue to hold. The simplest way to handle non-equality joins is to apply the predicates after all joins, while range predicates can be seen as a generalization of hash lookup: if a range predicate matches no results in a B+ tree index, we can safely backjump and delete just like  $\text{TTJ}$  does for hash lookup failures.

*Purpose of the no-good list.* We replaced Section 5 with the new  $\text{TTJ}^\epsilon$  variant which shows greater speedup compared to the optimizations that used to be in that section, and moved the no-good optimization to Appendix A. The benefit of the no-good list is data-dependent: when failures are concentrated on a few values, the no-good list will be small, and checking it can be faster than probing into a large hash table. Another case when the no-good list is useful is when the deletion happens after backjumping multiple levels. Then, checking the no-good list requires one single probe as opposed to probing into each level that was backjumped over.

*Relationship to Free Join.* We added a discussion of Free Join in Section 2. The new  $\text{TTJ}^\epsilon$  variant is directly inspired by Free Join, where we replace backjumping with early probing to avoid large intermediates.  $\text{TTJ}^\epsilon$  has stronger theoretical guarantees as compared to Free Join. Both  $\text{TTJ}$  and  $\text{TTJ}^\epsilon$  are instance-optimal for acyclic queries, while Free Join is not – the deletion step is crucial for instance-optimality.

## Response to Referee 2

Relevant changes in the paper are highlighted in **green**.

*Implementation in an existing system.* We appreciate how Yannakakis+ [34] brings the theoretical guarantee of Yannakakis' algorithm to existing systems, with a low barrier of adoption. However, we would like to argue that “ease of implementation” is a spectrum, and different design points in the space pose unique tradeoffs. Integrating  $\text{TTJ}$  (and the new  $\text{TTJ}^\epsilon$  variant) requires modifying the query execution engine, involving engineering effort that is out of the scope for the current paper. However, our experiments show this effort can be worthwhile, as our prototype implementation of  $\text{TTJ}^\epsilon$  significantly outperforms the rewrite-based Yannakakis+. Moreover, we believe the new  $\text{TTJ}^\epsilon$  variant makes implementation easier (compared to  $\text{TTJ}$ ) with the elimination of backjumping and deletion (the latter when possible). The Robust Predicate Transfer paper [40] presents yet another design point: their implementation using bloom filters is simpler than  $\text{TTJ}^\epsilon$  but still requires modifying the execution engine, and its performance sits between that of  $\text{TTJ}^\epsilon$  and Yannakakis+.

Nevertheless, the implementation of RPT requires changing hundreds of files, and has yet to be merged into DuckDB. Finally, we discuss challenges and potential solutions for integrating TTJ to SQLite in Appendix B.

*Left-deep plan and Bushy plans.* TTJ supports both left-deep and bushy plans, and we evaluate on bushy plans in Section 6.2. When the plan is bushy, we simply materialize each left-deep segment, as is commonly done for pipelined hash joins. We separate left-deep plans from bushy plans for theoretical analysis: TTJ guarantees to run in  $O(|IN| + |OUT|)$  only for left-deep plans and for each left-deep segment in a bushy plan, while the zero-overhead (in terms of hash probes) guarantee holds for arbitrary plans.

*Relationship with CROWN and Dynamic Yannakakis.* Thank you for pointing us to these papers, we added discussion on how TTJ relates to them in Section 2. TTJ is not the same as CROWN or Dynamic Yannakakis; in fact, it can be seen as a dual to these algorithms: while CROWN/Dynamic Yannakakis executes by traversing a join tree bottom-up, TTJ follows the join tree top-down (note that the bottom relation in a left-deep plan becomes the top of the join tree). As the referee pointed out, CROWN requires building the hash tables following a bottom-up order of the join tree, performing semijoins as it goes up. In contrast, TTJ imposes no constraint on the order of hash building, and they can even be built in parallel at the same time, or be pre-built as indices. This is because TTJ removes dangling tuples while probing the hash tables, whereas CROWN does so during the construction of the hash tables. Moreover, while TTJ guarantees to incur no additional hash table operations as compared to binary hash join on the same plan, bottom-up algorithms like CROWN cannot make such guarantees. Consider the case where the database contains no dangling tuples. Then, consider building the hash table for relation  $R_p$  with a child  $R_c$  in the join tree. In order to semijoin  $R_p$  with  $R_c$  on-the-fly during the hash table build, the algorithm must probe into  $R_c$  for each tuple in  $R_p$ . But because there are no dangling tuples, these probes will always return true. After all hash tables are built, to output the final join result the algorithm must perform the same probe again, and overall incurs 2x hash probes as compared to binary hash join.

*Definition of Tree Convolution.* Addressed above under common concerns.

*Comparison with Recent Work.* Addressed above under common concerns.

### Response to Referee 3

Relevant changes in the paper are highlighted in **purple**.

*Bloom and Bitmap Filters.* As addressed above under common concerns, we added a comparison with Robust Predicate Transfer [40] which is the most recent work using approximate filters for robust query processing. We also took inspiration from this line of work and designed a new variant (described in Section 5),  $TTJ^\epsilon$ , that performs better in practice. As our results show, the use of filters incurs additional overhead, while both TTJ and  $TTJ^\epsilon$  are provably zero-overhead when hash table operations are taken to be the main cost. We added a discussion of the related work in Section 2.

*Implementation of tuple deletion.* We use hash tables provided by the standard library. As we discuss in a footnote in Section 4.2, deletion is fast because we save a pointer to the tuple being deleted on the stack to avoid an additional hash table lookup, and each deletion occurs immediately after a hash probe failure, exhibiting good temporal locality. In the new  $TTJ^\epsilon$  variant, we further avoid deletion by taking advantage of the query structure as well as primary/foreign key constraints. We also discuss potential challenges surrounding deletion in the multi-threaded setting in Section 8.

*Tree convolution and Tree Decomposition.* Addressed above under common concerns.