

Low Latency Query Planning and Processing

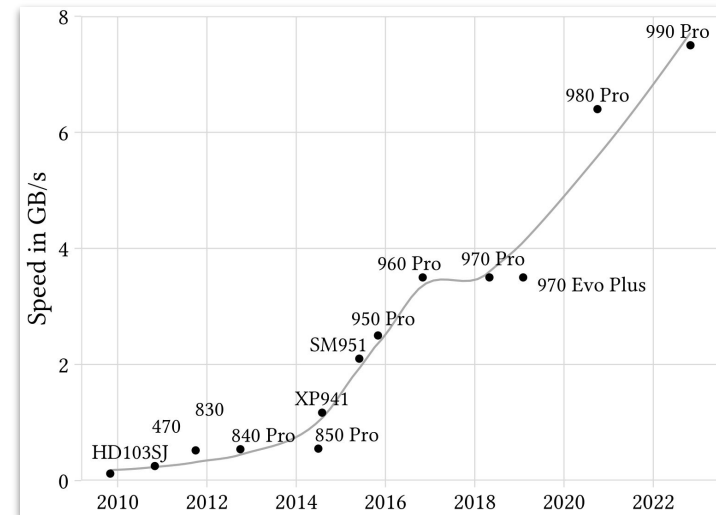
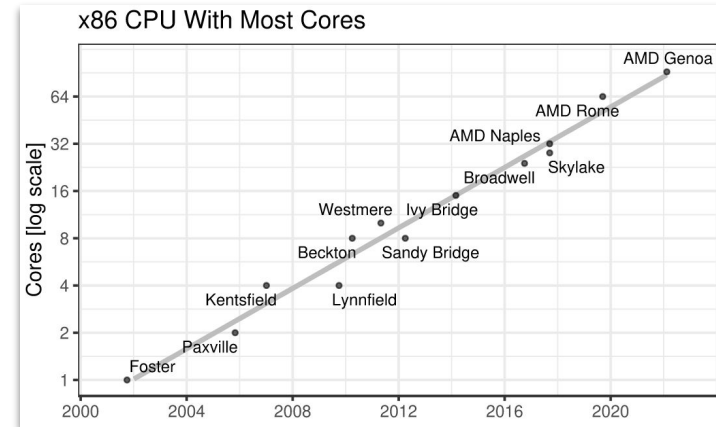
Philipp Fent

Hardware gets fast

- Large main memory
- Fast SSDs
- Many core machines

Low latency queries

- Still bound by CPU capabilities
- Algorithmic changes



Algorithmic challenges

Query processing

- Intra-query parallelism
- Shared state


Query planning

Cardinality estimation Algebra optimization

A Practical Approach to Groupjoin and Nested Aggregates

Philipp Fent
Technische Universität München
fentph@tum.de

Thomas Neumann
Technische Universität München
neumann@tum.de



ABSTRACT

Groupjoins, the combined execution of a join and a subsequent group-by, are common in analytical queries, and occur about 15% of the queries in TPC-H and TPC-DS. While they were originally invented to improve performance, efficient parallel execution of groupjoins can be limited by contention, which limits their usefulness in a many-core system. Having an efficient implementation of groupjoins is highly desirable, as groupjoins are not only used to fuse group-by and join but are also introduced by the unrolling components of the query optimizer to avoid nested-loop evaluation of aggregates. Furthermore, the query optimizer needs to be able to reason over the result of aggregation in order to schedule it correctly. Traditional selectivity and cardinality estimations quickly reach their limits when faced with complex relations from nested aggregates, which leads to poor cost estimations and thus, suboptimal query plans.

In this paper, we present techniques to efficiently estimate, plan, and execute estimation and nested aggregates. The proposed techniques to predict the result distrib-

tion of groupjoins are based on a novel filtering system that significantly helps the evaluation of groupjoins, which are up to a factor of 2

more efficient than the state-of-the-art.

ACM Reference Format: Fent, P., Neumann, T., and Neumann, T. 2022. A Practical Approach to Groupjoin and Nested Aggregates. In Proceedings of the ACM SIGMOD Conference on Management of Data, June 2022, 1–15.

KEYWORDS: Groupjoin, Nested Aggregates, Query Optimization, Cardinality Estimation, Algebra Optimization

The VLDB Journal (2022) 32:1:165–180
https://doi.org/10.1007/s00773-022-01045-x

SPECIAL ISSUE PAPER

Practical planning and execution of groupjoin and nested aggregates

Philipp Fent¹, Altan Birler¹, Thomas Neumann¹

Received: 1 April 2022 / Revised: 14 September 2022 / Accepted: 23 September 2022 / Published online: 22 October 2022
© The Author(s) 2022

Abstract

Groupjoins combine execution of a join and a subsequent group-by. They are common in analytical queries and occur in about 15% of the queries in TPC-H and TPC-DS. While they were originally invented to improve performance, efficient parallel execution of groupjoins can be limited by contention in many-core systems. Efficient implementations of groupjoins are highly desirable, as groupjoins are not only used to fuse group-by and join, but are also useful to efficiently execute nested aggregates. For these, the query optimizer needs to reason over the result of aggregation to optimally schedule it. Traditional systems quickly reach their limits of selectivity and cardinality estimations over complex columns and often treat group-by as an optimization barrier. In this paper, we present techniques to efficiently estimate, plan, and execute groupjoins and nested aggregates. We propose four novel techniques, aggregate estimators to predict the result distributions of aggregates, parallel groupjoin execution for scalable execution of groupjoins, index groupjoins, and a greedy query aggregation optimization technique that introduces nested aggregation to significantly improve execution plans. The resulting system has improved estimates, better execution plans, and a contention-free evaluation of groupjoins, which speeds up TPC-H and TPC-DS queries significantly.

Keywords: Query optimization · Query processing · Parallel processing

1 Introduction

Joins and aggregations are the backbone of query engines. A common query pattern, which we observe in many benchmarks [10,59] and industry applications [7], is a join with group aggregation on the same key:

```
SELECT cust_id, COUNT(*) SUM(r.value)
FROM customer cust, sales s
WHERE cust_id = s.cust_id
GROUP BY cust_id
```

In a traditional implementation, we answer the query by building two hash tables on the same key, one for the left join and one for the right join. However, we can speed up this query by reusing the join's hash table to also store the aggregate values. This combined execution of join and group-by is called a groupjoin [43].

The primary reason to use a groupjoin is its performance. A common query pattern, which we observe in many benchmarks [10,59] and industry applications [7], is a join with grouped aggregation on the same key:

```
SELECT cust_id, cnt, s
FROM customer cust, s
WHERE COUNT(*) AS cnt, SUM(r.value) AS s
FROM sales s
WHERE cust_id = s.cust_id
GROUP BY cust_id
```

Here, the query calculates a COUNT(*) over the inner table, which evaluates to zero when there are no join partners. Answering that query without nested-loop evaluation of the inner query is tricky, as a regular join directly group-by will produce wrong results for empty subqueries, which is known as the COUNT bug [58]. A groupjoin directly supports such

backbone of query engines. A feature in many benchmarks [10,59] is a join with grouped aggregation on the same key:

```
SELECT cust_id, cnt, s
FROM customer cust, s
WHERE COUNT(*) AS cnt, SUM(r.value) AS s
FROM sales s
WHERE cust_id = s.cust_id
```

A Practical Approach to Groupjoin and Nested Aggregates

Philipp Fent
Technische Universität München
fentph@tum.de

Thomas Neumann
Technische Universität München
neumann@tum.de

ABSTRACT

Groupjoins are essential for the efficient execution of queries. The necessary analysis, if we can and should apply optimizations and transform the query plan, is already challenging. Traditional techniques focus on the scalability of column or individual operators, which does not scale for analysis of data flow through the query. To make analysis of data flow through the query possible, which results in multi-second optimization time for deep algebra trees. Instead, we need to do the same for the entire algebra representation to efficiently support data flow analysis.

In this paper, we introduce Indexing Algebra, a novel representation of relational algebra that makes common optimization tasks efficient. Indexing Algebra enables efficient reasoning with an auxiliary index structure based on link-cut trees that support dynamic queries and queries in O(log n). This approach not only improves the complexity, but also allows adapted and concise filter mutations for the data flow questions needed for query optimization. While large queries are theoretically unbounded improvements, Indexing Algebra also improves optimization time of the relatively benchmark queries of TPC-H and TPC-DS by more than 12%.

ACM Reference Format: Fent, P., Birler, A., and Neumann, T. 2022. Asymptotically Better Query Optimization Using Indexed Algebra. In Proceedings of the ACM SIGMOD Conference on Management of Data, June 2022, 1–15.

KEYWORDS: Query Optimization, Algebra Optimization, Indexing Algebra, Link-Cut Trees, Query Optimization

INTRODUCTION

Optimizing the algebra plan of a query can take a significant portion of the overall runtime. The challenge for the optimizer here is that the data flow through the query, and its analysis, can be increasingly complex. Additionally, automatically generated queries with complex business logic simplify this problem [5, 18, 37]. Query optimizers struggle to deal with such complex queries, which is especially painful for small datasets where query optimization can be more expensive than query execution. Small data sizes are common during testing, but also in the real world, where, for example, Tableau reports that many workbooks contain fewer than a million rows [14]. As a result, query optimization usually operates on a budget, trading off optimization versus optimization time.

The work is funded under the Creative Commons BY-NC-ND 4.0 International license. For more information on this license, please visit <https://creativecommons.org/licenses/by-nc-nd/4.0/>. This work is part of the research project "Algebraic Query Optimization" funded by the German Research Foundation (DFG) under the special collaborative program SFB 1023/B1.

Proceedings of the VLDB Journal, Vol. 32, No. 1, 2022, 165–180.
doi:10.1007/s00773-022-01045-x

The primary reason to use a groupjoin, is its performance. We spend less time building hash tables, use less memory, and improve the responsiveness of this query. However, groupjoins are also more capable than regular group-by, as we can reuse the group explicitly. Consider the following nested query with subtly different semantics:

```
SELECT cust_id, cnt, s
FROM customer cust, s
WHERE COUNT(*) AS cnt, SUM(r.value) AS s
FROM sales s
WHERE cust_id = s.cust_id
```

A Practical Approach to Groupjoin and Nested Aggregates

Philipp Fent
Technische Universität München
fentph@tum.de

Guido Moerkotte
Technische Universität München
moerkotte@tum.de

ABSTRACT

Query optimization is essential for the efficient execution of queries. The necessary analysis, if we can and should apply optimizations and transform the query plan, is already challenging. Traditional techniques focus on the scalability of column or individual operators, which does not scale for analysis of data flow through the query. To make analysis of data flow through the query possible, which results in multi-second optimization time for deep algebra trees. Instead, we need to do the same for the entire algebra representation to efficiently support data flow analysis.

In this paper, we introduce Indexing Algebra, a novel representation of relational algebra that makes common optimization tasks efficient. Indexing Algebra enables efficient reasoning with an auxiliary index structure based on link-cut trees that support dynamic queries and queries in O(log n). This approach not only improves the complexity, but also allows adapted and concise filter mutations for the data flow questions needed for query optimization. While large queries are theoretically unbounded improvements, Indexing Algebra also improves optimization time of the relatively benchmark queries of TPC-H and TPC-DS by more than 12%.

ACM Reference Format: Fent, P., Birler, A., and Neumann, T. 2022. Asymptotically Better Query Optimization Using Indexed Algebra. In Proceedings of the ACM SIGMOD Conference on Management of Data, June 2022, 1–15.

KEYWORDS: Query Optimization, Algebra Optimization, Indexing Algebra, Link-Cut Trees, Query Optimization

INTRODUCTION

Optimizing the algebra plan of a query can take a significant portion of the overall runtime. The challenge for the optimizer here is that the data flow through the query, and its analysis, can be increasingly complex. Additionally, automatically generated queries with complex business logic simplify this problem [5, 18, 37]. Query optimizers struggle to deal with such complex queries, which is especially painful for small datasets where query optimization can be more expensive than query execution. Small data sizes are common during testing, but also in the real world, where, for example, Tableau reports that many workbooks contain fewer than a million rows [14]. As a result, query optimization usually operates on a budget, trading off optimization versus optimization time.

The work is funded under the Creative Commons BY-NC-ND 4.0 International license. For more information on this license, please visit <https://creativecommons.org/licenses/by-nc-nd/4.0/>. This work is part of the research project "Algebraic Query Optimization" funded by the German Research Foundation (DFG) under the special collaborative program SFB 1023/B1.

Proceedings of the VLDB Journal, Vol. 32, No. 1, 2022, 165–180.
doi:10.1007/s00773-022-01045-x

Asymptotically Better Query Optimization Using Indexed Algebra

Philipp Fent¹, Altan Birler¹, Thomas Neumann¹

Received: 1 April 2022 / Revised: 14 September 2022 / Accepted: 23 September 2022 / Published online: 22 October 2022
© The Author(s) 2022

Abstract

Groupjoins combine execution of a join and a subsequent group-by. They are common in analytical queries and occur in about 15% of the queries in TPC-H and TPC-DS. While they were originally invented to improve performance, efficient parallel execution of groupjoins can be limited by contention in many-core systems. Efficient implementations of groupjoins are highly desirable, as groupjoins are not only used to fuse group-by and join, but are also useful to efficiently execute nested aggregates. For these, the query optimizer needs to reason over the result of aggregation to optimally schedule it. Traditional systems quickly reach their limits of selectivity and cardinality estimations over complex columns and often treat group-by as an optimization barrier. In this paper, we present techniques to efficiently estimate, plan, and execute groupjoins and nested aggregates. We propose four novel techniques, aggregate estimators to predict the result distributions of aggregates, parallel groupjoin execution for scalable execution of groupjoins, index groupjoins, and a greedy query aggregation optimization technique that introduces nested aggregation to significantly improve execution plans. The resulting system has improved estimates, better execution plans, and a contention-free evaluation of groupjoins, which speeds up TPC-H and TPC-DS queries significantly.

Keywords: Query optimization · Query processing · Parallel processing

1 Introduction

Joins and aggregations are the backbone of query engines. A common query pattern, which we observe in many benchmarks [10,59] and industry applications [7], is a join with group aggregation on the same key:

```
SELECT cust_id, COUNT(*) SUM(r.value)
FROM customer cust, sales s
WHERE cust_id = s.cust_id
GROUP BY cust_id
```

In a traditional implementation, we answer the query by building two hash tables on the same key, one for the left join and one for the right join. However, we can speed up this query by reusing the join's hash table to also store the aggregate values. This combined execution of join and group-by is called a groupjoin [43].

Here, the query calculates a COUNT(*) over the inner table, which evaluates to zero when there are no join partners. Answering that query without nested-loop evaluation of the inner query is tricky, as a regular join directly group-by will produce wrong results for empty subqueries, which is known as the COUNT bug [58]. A groupjoin directly supports such



The work is funded under the Creative Commons BY-NC-ND 4.0 International license. For more information on this license, please visit <https://creativecommons.org/licenses/by-nc-nd/4.0/>. This work is part of the research project "Algebraic Query Optimization" funded by the German Research Foundation (DFG) under the special collaborative program SFB 1023/B1.

Proceedings of the VLDB Journal, Vol. 32, No. 1, 2022, 165–180.
doi:10.1007/s00773-022-01045-x



Groupjoin – Idea

- Combined execution of compatible join and aggregation
- Q: “Total sales per customer”



```
join = Hashtable()
for c in customer:
    join[c.id] = c
```

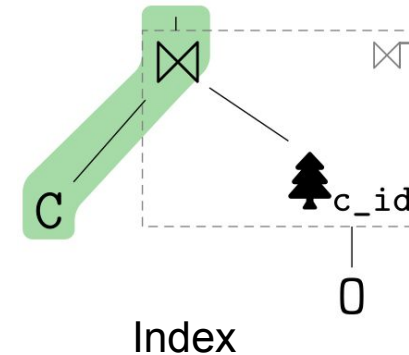

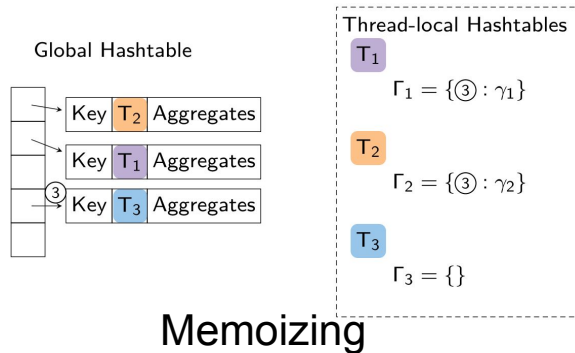

```
group = Hashtable()
for o in orders:
    if join.contains(o.c_id):
        group[c_id].sum += o.price
```

```
groupjoin = Hashtable()
for c in customer:
    groupjoin[c.id] = c

for o in orders:
    if groupjoin.contains(o.c_id):
        groupjoin[c_id].sum += o.price
```

Groupjoin – Avoiding contention

- Shared hash table unsuitable for multithreaded execution
- Four execution strategies for parallel groupjoin:



Nested Aggregates

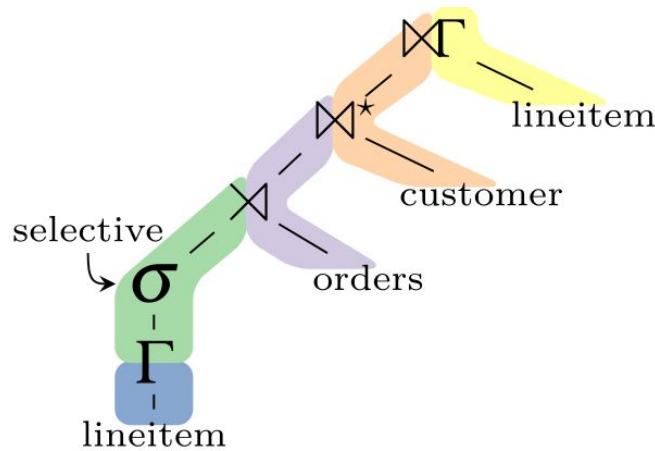
- Common in analytical queries
- HAVING predicates are hard to estimate

Q: “Large orders”


```
select l_orderkey
from lineitem
group by l_orderkey
having sum(l_quantity) > 300
```

Nested Aggregates

- Common in analytical queries
- HAVING predicates are hard to estimate
- But have significant impact on execution plans



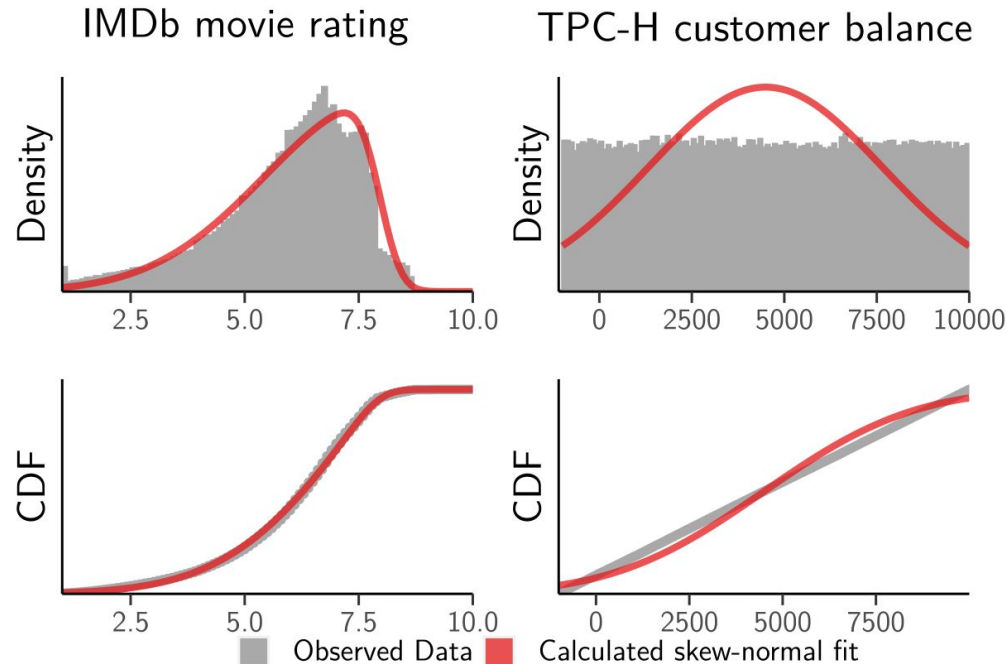
a) Selective σ -Predicate



b) Unselective σ -Predicate

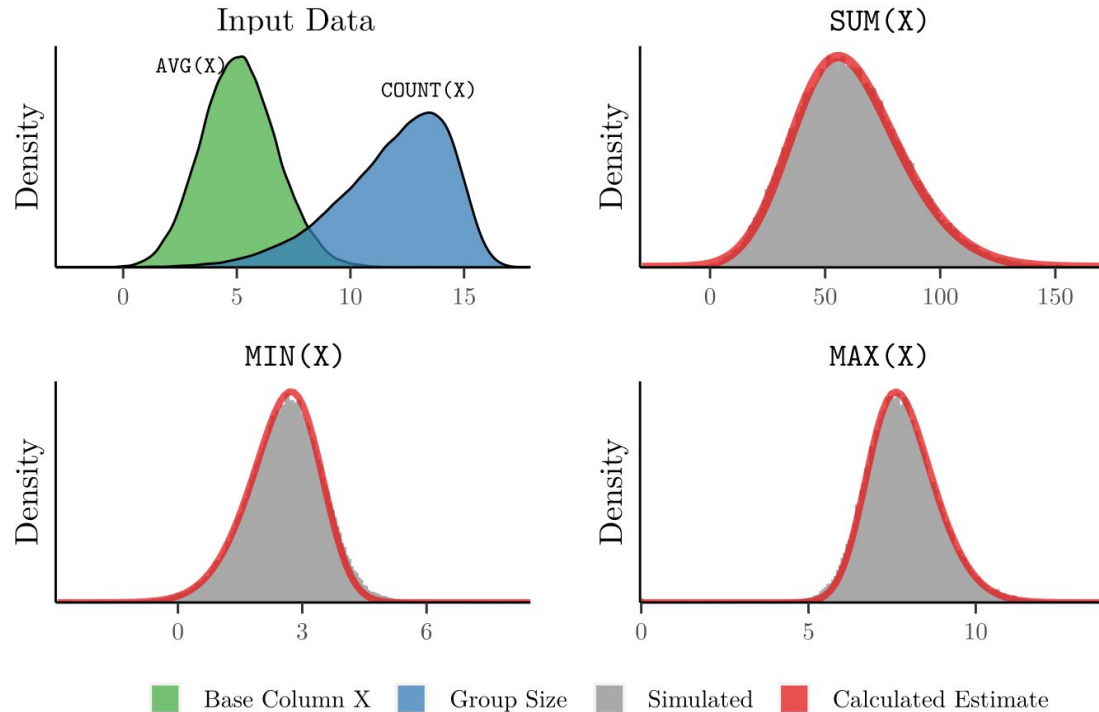
Estimating Aggregates

- Numerical columns $\sim \mathbf{N}(\mu, \sigma^2)$
- Cheap and generalizes nicely, but inherently symmetric



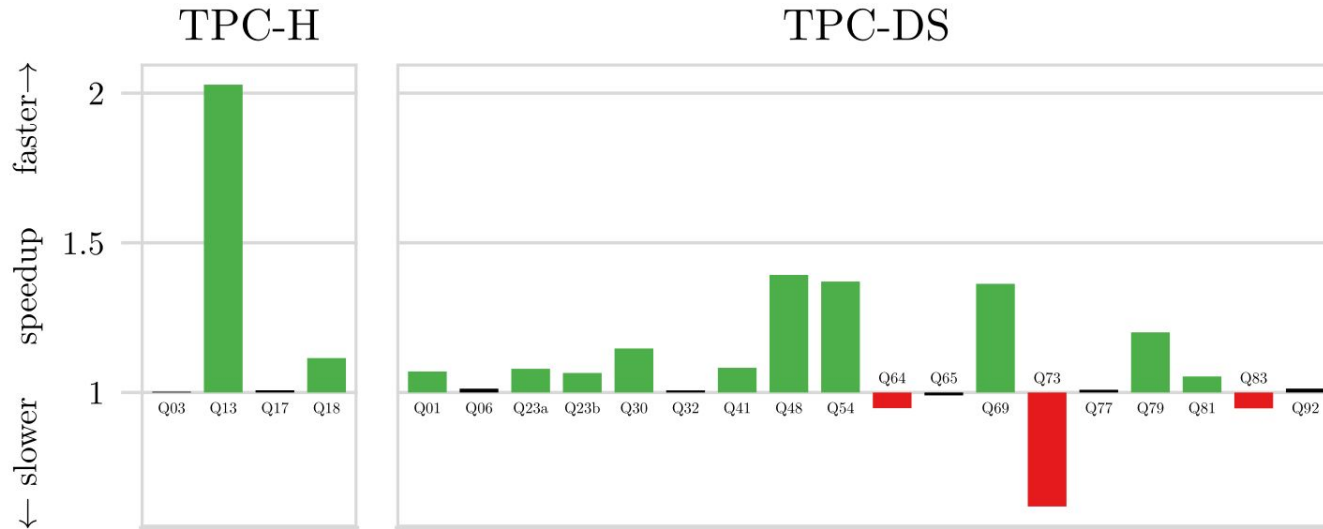
Estimating Aggregates

- Using a skew-normal distribution



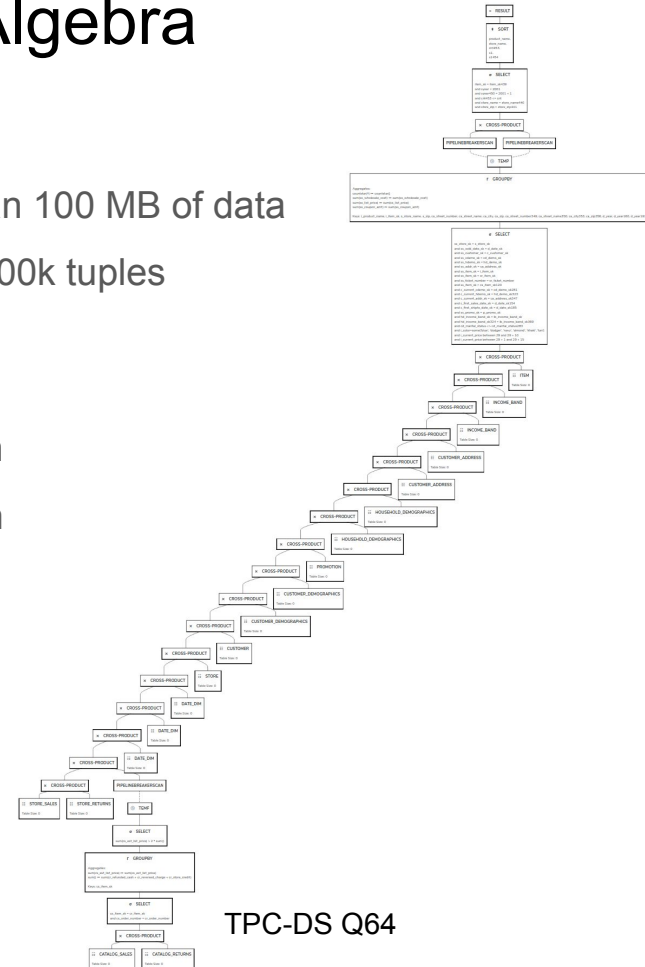
Practical Groupjoins and Nested Aggregates

- Effects $\frac{1}{8}$ of queries
- +23% in TPC-H, +6% in TPC-DS




Query Optimization with Indexed Algebra

- Complex queries on small workloads
 - BigQuery: 90% of queries processed less than 100 MB of data
 - Tableau Public: 90% of workbooks are less than 100k tuples
- TPC-H
 - Scale 1: 0.8 ms optimization, 20 ms execution
 - Scale 0.01: 0.8 ms optimization, 0.2 ms execution
- Optimization time scales super-linear with query complexity




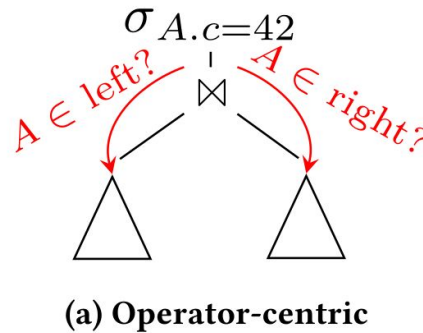
Algebra

- Relational algebra trees
 - Operators
 - Expressions
 - Columns / IUs
- Analyze data-flow for optimization
 - Which path?
 - Modifications?
 - Materialized?




Optimization

- Reason about the algebra to derive optimization possibilities
- Top-down, operator at a time
 - Needs $O(n^2)$ column sets
- Path-centric
 - Still $O(n)$ length
 - With indexing: $O(\log n)$




Indexing Algebra


- Index paths through the algebra
 - ➔ Faster path traversal
- Binary search trees on path depth
- Paths from root overlap
- **Link/cut trees** support that efficiently



(a) Represented algebra plan



(b) Balanced binary index of the path from B^6 to the root



(c) Index from D^5 to the root

Indexing Algebra


- Index paths through the algebra
 - ➔ Faster path traversal
- Binary search trees
on path depth
- Paths from root might overlap
- **Link/cut trees** support that efficiently



Rel. Algebra	Transformation	Traversal
w/o index	$O(1)$	$O(n)$
static index	$O(n)$	$O(\log n)$
path labeling	$O(n)$	$O(1)$
Indexed Algebra	$O(\log n)$	$O(\log n)$

Indexed Algebra Performance

- Significant overall improvements
- 10 - 30% faster optimization
- 8% better *end-to-end* latency in Tableau Public



Conclusion

Query processing

✓ Intra-query parallelism

✓ Shared state

Query planning


✓ Cardinality estimation

✓ Algebra operations

A Practical Approach to Groupjoin and Nested Aggregates

Philipp Fent
Technische Universität München
fentph@tum.de

Thomas Neumann
Technische Universität München
neumann@tum.de



ABSTRACT
Groupjoins, the combined execution of a join and a subsequent group-by, are common in analytical queries, and occur in about 15% of the queries in TPC-H and TPC-DS. While they were originally invented to improve performance, efficient parallel execution of groupjoins can be limited by contention, which limits their usefulness in a many-core system. Having an efficient implementation of groupjoins is highly desirable, as groupjoins are not only used to fuse group-by and join but are also introduced by the unrolling of aggregates. Furthermore, the query optimizer needs to be able to reason over the result of aggregation in order to schedule it correctly. Traditional selectivity and cardinality estimations quickly reach their limits when faced with complex relations from nested aggregates, which leads to poor cost estimations and thus, suboptimal query plans.

In this paper, we present techniques to efficiently estimate, plan, and execute estimation and nested aggregates. The proposed two-phase approach to predict the result distribution of groupjoins uses a scalable filtering system that significantly helps the evaluation of groupjoins, which are up to a factor of 2.5 faster.

KEYWORDS
Groupjoin, Nested Aggregates, Query Optimization, Parallel Processing

The VLDB Journal (2023) 32:1165–1180
https://doi.org/10.1007/s00778-022-00540-x

SPECIAL ISSUE PAPER

Practical planning and execution of groupjoin and nested aggregates

Philipp Fent¹ · Altan Birler¹ · Thomas Neumann¹

Received: 1 April 2022 / Revised: 14 September 2022 / Accepted: 23 September 2022 / Published online: 22 October 2022
© The Author(s) 2022

Abstract
Groupjoins combine execution of a join and a subsequent group-by. They are common in analytical queries and occur in about 15% of the queries in TPC-H and TPC-DS. While they were originally invented to improve performance, efficient parallel execution of groupjoins can be limited by contention in many-core systems. Efficient implementations of groupjoins are highly desirable, as groupjoins are not only used to fuse group-by and join, but are also useful to efficiently execute nested aggregates. For these, the query optimizer needs to reason over the result of aggregation to optimally schedule it. Traditional systems quickly reach their limits of selectivity and cardinality estimations over complex relations and often treat group-by as an optimization barrier. In this paper, we present techniques to efficiently estimate, plan, and execute groupjoins and nested aggregates. We propose four novel techniques, *aggregate estimators* to predict the result distributions of aggregates, *parallel groupjoin execution* for scalable execution of groupjoins, *index groupjoins*, and a *greedy query aggregation optimization* technique that introduces nested aggregations to significantly improve execution plans. The resulting system has improved estimates, better execution plans, and a contention-free evaluation of groupjoins, which speeds up TPC-H and TPC-DS queries significantly.

Keywords Query optimization · Query processing · Parallel processing

1 Introduction

Joins and aggregations are the backbone of query engines. A common query pattern, which we observe in many benchmarks [10,59] and industry applications [7], is a join with group aggregation on the same key:

```
SELECT cust_id, COUNT(*) SUM(n.value)
FROM customer cust, sales s
WHERE cust_id = s.cust_id
GROUP BY cust_id
```

In a traditional implementation, we answer the query by building two hash tables by the same key, one for the left join and one for the group-by. However, we can speed up this query by reusing the join's hash table to also store the aggregate values. This combined execution of join and group-by is called a *groupjoin* [43].

The primary reason to use a groupjoin is its performance. We spend less time building hash tables, use less memory, and improve the responsiveness of this query. However, groupjoins are also more capable than regular group-bys, as we can create the group explicitly. Consider the following nested query, with subtly different semantics:

```
SELECT cust_id, cnt, s
FROM customer cust, s
WHERE COUNT(*) AS cnt, SUM(n.value) AS s
```

Here, nested query calculates a `COUNT(*)` over the inner table, which evaluates to zero when there are no join partners. Answering that query without nested-join evaluation of the inner query is tricky, as a regular join directly group-by will produce wrong results for empty subqueries, which is known as the `COUNT(*) bug` [58]. A groupjoin properly supports such

A Practical Approach to Groupjoin

Philipp Fent
Technische Universität München
fentph@tum.de

Guido Moerkotte
Technische Universität München
moerkotte@tum.de

ABSTRACT
Query optimization is essential for the efficient execution of queries. The necessary analysis, if we can and should apply optimizations and transform the query plan, is already challenging. Traditional techniques focus on the availability of operators as individual objects, which does not scale for analysis of data flow through the query. To make analysis more efficient, we propose table quadratic queries, which result in multi-recursive optimization flow for deep algebra trees. Instead, we need to do the same matrix algebra representation to efficiently support data flow analysis.

In this paper, we introduce Indexed Algebra, a novel representation of relational algebra that makes common optimization tasks efficient. Indexed Algebra enables efficient reasoning with an auxiliary index structure based on link-cut trees that support dynamic updates and queries in $O(\log n)$. This approach not only improves the complexity, but also allows adapted and concise flow models for the data flow questions needed for query optimization. While large queries are theoretically unsolvable improvements, Indexed Algebra also improves optimization time of the relatively harmless queries of TPC-H and TPC-DS by more than 12%.

KEYWORDS
Query optimization, Query processing, Parallel processing

Philipp Fent, Guido Moerkotte, and Thomas Neumann, *Asymptotically Better Query Optimization Using Indexed Algebra*, VLDB 2023, 1165–1180, 2023.

https://doi.org/10.1007/s00778-022-00540-x

VLDB Archive Availability:
The source code, data, and/or other artifacts have been made available at <https://github.com/infodiv/IndexedAlgebra>.

1 INTRODUCTION

Optimizing the algebra plan of a query can take a significant portion of the overall runtime. The challenge for the optimizer here is that the data flow through the query, and its analysis can be increasingly complex. Additionally, automatically generated queries with complex business logic simplify this problem [5, 18, 37]. Query optimizers struggle to deal with such complex queries, which is especially painful for small datasets where query optimization can be more expensive than query execution. Small data sizes are common during testing, but also in the real world, where, for example, Tableau reports that many workbooks contain fewer than a million table [19]. As a result, query optimization usually operates on a budget, trading off optimization versus optimization time.

The work is funded under the Creative Commons BY-NC-ND 4.0 International license. For more information on data privacy, please refer to our website at <https://www.lmu.de/en/research/infodiv>. Copyright is held by the respective publisher. Publication rights reserved by the VLDB Foundation.

Proceedings of the VLDB Journal, Vol. 36, No. 11, 2023, 1165–1180.
doi:10.1007/s00778-022-00540-x

The primary reason to use a groupjoin, is its performance. We spend less time building hash tables, use less memory, and improve the responsiveness of this query. However, groupjoins are also more capable than regular group-bys, as we can create the group explicitly. Consider the following nested query with subtly different semantics:

```
SELECT cust_id, cnt, s
FROM customer cust, s
WHERE COUNT(*) AS cnt, SUM(n.value) AS s
FROM sales s, l, s, L, L, L
```

Asymptotically Better Query Optimization Using Indexed Algebra

Philipp Fent
Technische Universität München
fentph@tum.de

Guido Moerkotte
Technische Universität München
moerkotte@tum.de

ABSTRACT
Query optimization is essential for the efficient execution of queries. The necessary analysis, if we can and should apply optimizations and transform the query plan, is already challenging. Traditional techniques focus on the availability of operators as individual objects, which does not scale for analysis of data flow through the query. To make analysis more efficient, we propose table quadratic queries, which result in multi-recursive optimization flow for deep algebra trees. Instead, we need to do the same matrix algebra representation to efficiently support data flow analysis.

In this paper, we introduce Indexed Algebra, a novel representation of relational algebra that makes common optimization tasks efficient. Indexed Algebra enables efficient reasoning with an auxiliary index structure based on link-cut trees that support dynamic updates and queries in $O(\log n)$. This approach not only improves the complexity, but also allows adapted and concise flow models for the data flow questions needed for query optimization. While large queries are theoretically unsolvable improvements, Indexed Algebra also improves optimization time of the relatively harmless queries of TPC-H and TPC-DS by more than 12%.

KEYWORDS
Query optimization, Query processing, Parallel processing

Philipp Fent, Guido Moerkotte, and Thomas Neumann, *Asymptotically Better Query Optimization Using Indexed Algebra*, VLDB 2023, 1165–1180, 2023.

https://doi.org/10.1007/s00778-022-00540-x

VLDB Archive Availability:
The source code, data, and/or other artifacts have been made available at <https://github.com/infodiv/IndexedAlgebra>.

1 INTRODUCTION

Optimizing the algebra plan of a query can take a significant portion of the overall runtime. The challenge for the optimizer here is that the data flow through the query, and its analysis can be increasingly complex. Additionally, automatically generated queries with complex business logic simplify this problem [5, 18, 37]. Query optimizers struggle to deal with such complex queries, which is especially painful for small datasets where query optimization can be more expensive than query execution. Small data sizes are common during testing, but also in the real world, where, for example, Tableau reports that many workbooks contain fewer than a million table [19]. As a result, query optimization usually operates on a budget, trading off optimization versus optimization time.

The work is funded under the Creative Commons BY-NC-ND 4.0 International license. For more information on data privacy, please refer to our website at <https://www.lmu.de/en/research/infodiv>. Copyright is held by the respective publisher. Publication rights reserved by the VLDB Foundation.

Proceedings of the VLDB Journal, Vol. 36, No. 11, 2023, 1165–1180.
doi:10.1007/s00778-022-00540-x




Fig. 1: Relation algebra tree with subtle data flow. In this paper, we optimize queries by efficiently analyzing data flow.

Some typical questions that come up during query optimization are from which part of the plan does a value come from? What are the join partners? Can we push a predicate down into the input? Consider for example the SQL query below:

```
SELECT *
FROM A, B, C LEFT OUTER JOIN D ON C.U = D.U
WHERE A.X = T AND A.X = B.X AND B.X = C.X
AND C.V = T AND D.Z = B
```

In this small example it is easy to see which attributes form join edges (e.g., `SQL`), which filters can be pushed down (e.g., `WHERE`) and not directly (e.g., `WHERE`). In general, these questions are difficult because the FROM clause can contain arbitrary subqueries. The traditional solution to this problem is to keep track of the columns that are available in each step of the query in a set [5, 16] and to more proactively answer queries by step, checking the available columns for each transformation. But if we have a join tree of depth n , where each join produces at least one column, the construction time for these column sets grows with $O(n^n)$, which is highly unattractive for large queries.

Even if we ignore the performance problems, this naive loop of individual operators is insufficient to express optimizations efficiently. In many cases we want to inspect the full data flow instead. Consider the small algebra tree shown in Figure 1. The top most join produces in this case computes an attribute from its left input ($A.V$) with an attribute from its right input ($C.V$). While this data flow direction might be easy to see for such a small example, it is non-obvious when there are dozens of operators between the base tables and the predicate. And note that the example tree contains a non-trivial data flow that is not obvious as a first glance: The selection operator on the lower right uses an attribute that is produced in a different part of the operator tree, which effectively makes the top-most operator a dependent join. Evaluating such a join is highly non-trivial, since it requires a nested loop query execution. The query optimizer has to detect these dependent joins and can then rewrite the query to remove the correlation between parts of the join tree [17]. While we can detect these dependent joins by

© Philipp Fent
fentph@tum.de
Altan Birler
altan.birler@tum.de
Thomas Neumann
neumann@tum.de
1 Technische Universität München, Garching, Germany

