CERT: Finding Performance Issues in Database Systems Through the Lens of Cardinality Estimation

Jinsheng Ba
National University of Singapore
Singapore
bajinsheng@u.nus.edu

Manuel Rigger
National University of Singapore
Singapore
rigger@nus.edu.sg

ABSTRACT

Database Management Systems (DBMSs) process a given query by creating a query plan, which is subsequently executed, to compute the query's result. Deriving an efficient query plan is challenging, and both academia and industry have invested decades into researching query optimization. Despite this, DBMSs are prone to performance issues, where a DBMS produces an unexpectedly inefficient query plan that might lead to the slow execution of a query. Finding such issues is a longstanding problem and inherently difficult, because no ground truth information on an expected execution time exists. In this work, we propose *Cardinality Estimation Restriction Testing* (*CERT*), a novel technique that finds performance issues through the lens of cardinality estimation. Given a query on a database, *CERT* derives a more restrictive query (*e.g.*, by replacing a **LEFT JOIN** with an **INNER JOIN**), whose estimated number of rows should not exceed the estimated number of rows for the original query. *CERT* tests cardinality estimation specifically, because it was shown to be the most important part for query optimization; thus, we expect that finding and fixing cardinality-estimation issues might result in the highest performance gains. In addition, we found that other kinds of query optimization issues can be exposed by unexpected estimated cardinalities, which can also be found by *CERT*. *CERT* is a black-box technique that does not require access to the source code; DBMSs expose query plans via the **EXPLAIN** statement. *CERT* eschews executing queries, which is costly and prone to performance fluctuations. We evaluated *CERT* on three widely used and mature DBMSs, MySQL, TiDB, and CockroachDB. *CERT* found 13 unique issues, of which 2 issues were fixed and 9 confirmed by the developers. We expect that this new angle on finding performance bugs will help DBMS developers in improving DMBSs' performance.

CCS CONCEPTS

• **Software and its engineering** → **Software testing and debugging**; **Software performance**.

KEYWORDS

Database, Performance Issue, Cardinality Estimation

1 INTRODUCTION

Database Management Systems (DBMSs) are fundamental software systems that allow users to retrieve, update, and manage data [25, 48, 56]. Most DBMSs support the *Structured Query Language* (SQL), which allows users to specify queries in a declarative way. Subsequently, the DBMSs translate the query into a concrete execution plan. To balance the trade-off between spending little time on optimization, which is performed at run time, and finding an efficient execution plan, researchers and practitioners have invested decades of effort into query optimization, covering directions such as search space exploration for join ordering [12, 13, 35], index data structures [16], execution time prediction [1, 57], or parallel execution on multi-core CPUs [15] and GPUs [37].

Finding performance issues in DBMSs—also referred to as optimization opportunities or performance bugs—is challenging. Given a query Q and a database D, we want to determine whether executing Q on D results in unexpectedly suboptimal performance. In general, no ground truth is available that specifies whether Q executes within a reasonable time. To exacerbate this issue, DBMSs use various heuristics and cost models during optimizations, or make trade-offs in optimizing specific kinds of queries over others. Second, the execution time of Q might be significant if D is large, making it time-consuming to measure Q's actual performance. Given that the execution time depends on various factors of the execution environment [32] (*e.g.*, the state of caches), it might even be necessary to execute Q multiple times to obtain a reasonably reliable measure of its execution time. Cloud environments are in particular prone to noise [24]; a report on testing SAP HANA [2] has recently stressed that performance testing for cloud offerings of DBMSs—such as SAP HANA Cloud, which runs in Kubernetes pods—is one of the main challenges in testing DBMSs due to inherently noisy environments.

Benchmark suites such as TPC-DS [51] or TPC-H [33] are widely used in practice to monitor DBMSs' performance over versions through predetermined performance baselines, which could be specified [41, 61, 62]. However, deriving an appropriate baseline is challenging and might result in false alarms. Automated testing techniques have been proposed to find performance issues without the need of curating a benchmark suite. *APOLLO* [21] generates databases and queries automatically and validates whether executing the query on different versions of the DBMS results in significantly different execution times. However, *APOLLO* can only find regression bugs. *AMOEBA* [31] finds performance issues by examining discrepancies in the execution time of a pair of semantically equivalent queries. However, semantically equivalent queries do not necessarily exhibit a similar performance as the issues found by *AMOEBA* have a high false positive rate—only 6 of 39 issues

were confirmed by the developers, 5 of which were fixed [31]. For the above methods, queries need to be executed on sufficiently large databases to detect significant performance discrepancies.

In this work, we propose *Cardinality Estimation Restriction Testing* (*CERT*), a general technique that finds performance issues by testing the DBMSs' cardinality estimation. *Cardinality estimation* is the process in which *cardinality estimator* computes *estimated cardinalities*, the estimated numbers of rows that will be returned. Since estimated cardinalities are approximate, it is infeasible to check for a specific number. Rather, the core idea of our approach is that making a given query more restrictive should cause the cardinality estimator to estimate that the more restrictive query should fetch at most as many rows as the original query. More formally, given a query Q and a database D, $Card(Q, D)$ denotes the actual cardinality, that is, the exact number of records to be fetched by Q on D. If we derive a more restrictive query Q' from Q, $Card(Q', D) \leq Card(Q, D)$ always holds. $EstCard(Q, D)$ denotes the estimated cardinality for Q, and we expect $EstCard(Q', D) \leq EstCard(Q, D)$ to also hold for any DBMS. We refer to this property as *cardinality restriction monotonicity*. Any violation of this property indicates a potential performance issue.

CERT addresses the aforementioned challenges. Cardinality estimation accuracy was shown to be the single most important component for deriving an efficient execution plan [26]. Therefore, we believe that pinpointing issues in cardinality estimation would help developers focus on the most relevant issues, addressing which might result in significant performance gains. Additionally, this idea is applicable to finding a broader range of performance issues. For example, we found that other kinds of query optimization issues can be exposed by unexpected estimated cardinalities, as shown in Listing 3. Estimated cardinalities can be readily obtained by DBMSs without executing Q; DBMSs typically provide a SQL **EXPLAIN** statement that provides this information as part of a query plan, allowing our technique to achieve high throughput. Furthermore, since our method does not measure run-time performance, it can be used in noisy environments, and minimal test cases that demonstrate the performance issue can be automatically obtained [65]. Finally, *CERT* is a black-box technique, because applying it requires no access to the source code or binary.

Listing 1 shows a running example demonstrating *CERT*. We randomly generate SQL statements as shown in lines 1–4 to create a database state and ensure that each table's data statistics are up to date in lines 5–6. Then, we randomly generate a query with a **LEFT JOIN** and derive a more restrictive query by replacing the **LEFT JOIN** with an **INNER JOIN** as shown in lines 8–9. The second query is more restrictive than the first query as **INNER JOIN** should always fetch no more rows than **LEFT JOIN**. We examined the estimated cardinalities in their query plans, obtained by using an **EXPLAIN** statement. Despite having made the query more restrictive, the cardinality estimator estimated that the first query fetches 20 rows, while the second one fetches 60 rows, which is unexpected. The root cause was an incorrect double-counting when estimating the selectivity of **OR** expression in the **ON** condition of the **INNER JOIN**. The estimated cardinality of the first query with **LEFT JOIN** should be no less than 60; after the developers fixed this issue, the estimated cardinality was changed to 60. This fix improved the performance of

Listing 1: This running example demonstrates a performance issue found by *CERT* in CockroachDB. It is due to an incorrect double-counting when estimating the selectivity of OR expressions in join ON conditions.

```
1   CREATE TABLE t0 (c0 INT);
2   CREATE TABLE t1 (c0 INT);
3   INSERT INTO t0 VALUES (1), (2), (3), (4), (5), (6), (7),
        (8), (9), (10), (11), (12), (13);
4   INSERT INTO t1 VALUES (21),(22),(23),(24),(25);
5   ANALYZE t0;
6   ANALYZE t1;
7
8   EXPLAIN SELECT * FROM t0 LEFT JOIN t1 ON t0.c0<1 OR
        t0.c0>1; -- estimated rows: 20 🐞
9   EXPLAIN SELECT * FROM t0 INNER JOIN t1 ON t0.c0<1 OR
        t0.c0>1; -- estimated rows: 60 ✔
10  --------------------------------------------------
11  • cross join(left outer)  • cross join
12  | estimated row:20        | estimated row:60
13  | pred:(c0<1)OR(c0>1)      |-• filter
14  |-• scan                   | | estimated row:12
15  |    estimated row:13      | | filter:(c0<1)OR(c0>1)
16  |    table: t0@t0_pkey     | |-• scan
17  |-• scan                   | |    estimated row:13
18  |    estimated row:5       | |    table: t0@t0_pkey
19  |    table: t1@t1_pkey     |-• scan
20                            |    estimated row:5
21                            |    table: t1@t1_pkey
```

the query **SELECT * FROM t0 LEFT OUTER JOIN t1 ON t0.c0<1 OR t0.c0>1 FULL JOIN t2 ON t0.c0=t2.c0** by 20% as shown in Listing 7. The improvement was due to a more accurate estimated cardinality, which enabled a better selection of the join order. Note that we avoided executing the query; *CERT* only examines query plans.

We implemented *CERT* in *SQLancer*, a popular DBMS testing tool, and evaluated it on three widely used and mature DBMSs, MySQL, TiDB, and CockroachDB. While MySQL is one of the most popular open-source DBMSs, TiDB and CockroachDB are developed by companies. We reported 14 performance issues to the developers, who confirmed that 13 of them were unique and 12 were unknown. Of these unique issues, 2 issues were fixed, 9 other issues were confirmed, and 2 issues required further investigation. Similar to existing work, *CERT* might report false alarms, since implementations might not strictly adhere to the *cardinality restriction monotonicity*. However, in practice, none of the issues that we reported were considered false alarms. Our evaluation demonstrates the high throughput achieved by eschewing executing queries; our implementation can validate 386× more queries than *AMOEBA* in the same time period. We believe that these results demonstrate that *CERT* might become a standard technique in DBMS developers' toolbox, due to its efficiency and effectiveness, and hope that it will inspire future work on finding performance issues in DBMSs.

Overall, we make the following contributions:

- We present a motivational study to investigate the causes of previous performance issues.
- We propose a novel technique, *CERT*, to test cardinality estimation for finding performance issues in query optimization

Listing 2: The EBNF representation of a query.

```
1  SELECT [DISTINCT]
2    select_expression (, select_expression)*
3    FROM table_reference (INNER | LEFT | RIGHT | FULL |
         CROSS JOIN table_reference)*
4    (WHERE predicate)+
5    (GROUP BY predicate
6    (HAVING predicate)+)+
7    (LIMIT row_count)+ ;
```

without measuring execution time. We show a concrete realization of the technique by proposing 12 query-restriction rules.

- We implemented *CERT* in *SQLancer* and evaluated it on multiple aspects. *CERT* found 13 unique issues of cardinality estimation in widely-used DBMSs, and 11 issues were confirmed or fixed. The source code of *CERT* is publicly available, and has been integrated into *SQLancer*.[1]

2 BACKGROUND

Structured Query Language. Structured Query Language (SQL) [5] is a declarative programming language that expresses only the logic of a computation without specifying its specific execution. SQL is widely supported by DBMSs; for example, according to a popular ranking,[2] the 10 most popular DBMSs support it. Listing 2 shows the EBNF representation [36], a metasyntax notation to express context-free grammars, of a SQL query, whose features we considered in this work. A query starts with the **SELECT** keyword. It can optionally be succeeded by a **DISTINCT** clause that specifies that only unique records should be returned. A **JOIN** clause joins two tables or views; various joins exist that differ on whether and what rows should be joined when the join predicate evaluates to false. A query can contain a single **WHERE** clause; only rows for which its predicate evaluates to true are included in the result set. Similar to **DISTINCT**, the **GROUP BY** clause groups rows that have the same values into a single row. It can be followed by a **HAVING** clause that excludes records after grouping them. The **LIMIT** clause is used to restrict the number of records that are fetched. More advanced features, such as window functions, common table expressions (CTEs), and subqueries can be used. While we did not consider them in this work, we believe that our proposed approach could be extended to support them.

Query optimization. DBMSs include query optimizers that, after parsing the SQL query, determine an efficient *query plan*, which specifies how a SQL query is executed. Determining an efficient query plan is challenging, since many factors might influence the plan's performance. The most commonly used models are cost-based [6]—the query plan with the lowest projected performance cost is chosen. Cardinality estimation was found to be the most important factor that affects the quality of query optimization [26]. Cardinality estimators typically obtain data statistics of the tables to be queried by sampling [18], through histograms [46], or machine learning algorithms [10, 23, 59]. Then, they enumerate all sub-plan queries, which are queries that process only a subset of tables in a

query, and estimate how many rows they fetch. For example, for a query $A \bowtie B \bowtie C$ ($\bowtie$ denotes a join), cardinality estimators could estimate the cardinalities of A, B, C respectively, and then estimate the cardinalities of $A \bowtie B$, $B \bowtie C$, $A \bowtie C$, and $A \bowtie B \bowtie C$. Lastly, the estimated cardinalities of these sub-plan queries help to decide the join order—whether $A \bowtie B$ or $B \bowtie C$ should be executed first.

Query plan. A query plan is a tree of operations that describes how a specific DBMS executes a SQL statement. A query plan can be obtained by executing a query with the prefix **EXPLAIN**. Query plans typically include estimated cardinalities for operations that affect the number of subsequent rows. Listing 1 shows the two query plans for the two queries of the running example in lines 11–21. The first query uses a **LEFT JOIN**, and its query plan includes three operations: **cross join**, **scan**, and **scan**. The second query uses an **INNER JOIN**, and its query plan includes four operations: **cross join**, **filter**, **scan**, and **scan**. The structural difference between both query plans is the location of operation **filter**, as the predicate (c0<1)**OR**(c0>1) in the **ON** clause of both queries is part of the **cross join** of the first query plan, but is a separate operation **filter** in the second query plan. For the first query plan, the estimated cardinality 20 of the root node **cross join** (**left outer**) is determined based on the predicate as well as the two estimates for the number of records in tables t0 and t1, which are 13 and 5. For the second query plan, the estimated cardinality for **cross join** is 60; here, the estimate is partly based on the **filter** operation, which is estimated to return 12 rows. In this work, when we refer to the estimated cardinality of a query, we refer to the query's root node.

3 PERFORMANCE ISSUE STUDY

As a motivating study, to investigate if performance issues are caused by incorrect cardinality estimation in practice, we examined previous performance issues related to query optimization.

Subjects. We studied the issues reported for MySQL, TiDB, and CockroachDB. MySQL is the most popular relational DBMS according to a survey in 2021.[3] TiDB and CockroachDB are popular enterprise-class DBMSs, and their open versions on GitHub are highly popular as they have been starred more than 33k and 26k times. They are widely used and have thus been studied in other DBMS testing works [30, 43, 44].

Methodology. We searched for performance issues using the keywords "*slow*" or "*suboptimal*" in the above-stated DBMSs, aiming to obtain issues that relate to either slow execution or suboptimal query plans. For MySQL, we chose issues whose status was *closed*, the severity was *(S5) performance*, and the type was *MySQL Server: Optimizer* in the bug tracker.[4] Considering MySQL was first released in 1995 and some issues are too old to be reproduced, we investigated the issues in version 5.5 or later. For TiDB, we searched its repository[5] by the filter *is:issue is:closed linked:pr label:type/bug slow in:title*. For CockroachDB, we searched its repository[6] by the filter *is:issue is:closed linked:pr label:C-bug slow in:title*. Then, we

---

Table 1: Previous performance issues.

DBMS	#ID	Caused by Cardinality Estimation
MySQL	61631	✓
MySQL	56714	✓
MySQL	25130	X
CockroachDB	93410	X
CockroachDB	71790	X
TiDB	9067	✓
Sum	6	3

Listing 3: The performance issue #56714 in MySQL.

```
1   CREATE TABLE test (a INT PRIMARY KEY AUTO_INCREMENT, b
        INT NOT NULL, INDEX (b)) engine=INNODB;
2   CREATE TABLE integers(i INT UNSIGNED NOT NULL);
3   INSERT INTO integers(i) VALUES (0), (1), (2), (3), (4),
        (5), (6), (7), (8), (9);
4   INSERT INTO test (b)
5   SELECT units.i MOD 2
6   FROM integers AS units
7       CROSS JOIN integers AS tens
8       CROSS JOIN integers AS hundreds
9       CROSS JOIN integers AS thousands
10      CROSS JOIN integers AS tenthousands
11      CROSS JOIN integers AS hundredthousands;
12
13  EXPLAIN SELECT MAX(a) FROM test WHERE b=0; -- estimated
        rows: {500360} 🐛, {1} ✔
```

manually analyzed and reproduced each issue to identify whether it is a performance issue related to query optimization and caused by cardinality estimation. Specifically, if the estimated cardinality of the query in a report was changed by the fix, we deemed the performance issue to be caused by incorrect cardinality estimation.

Analysis. Table 1 shows the studied performance issues. Overall, we identified six issues in three DBMSs, and three of them were caused by incorrect cardinality estimation. We attribute the lower number of performance issues to the difficulty in identifying and resolving performance issues in query optimization. For issues #61631 and #56714, they produce inefficient query plans that have higher estimated cardinalities than the optimal query plans. Although both issues are not directly due to the faults in cardinality estimators, they are still observable through estimated cardinalities, so we deemed both were caused by cardinality estimation. Issue #9067 was caused by the cardinality estimation due to an issue in calculating cardinality for correlated columns. The other three issues, which were not caused by cardinality estimation, were due to inefficient operations. For example, issue #71790 was due to the inefficient implementation of **MERGE JOIN**, which did not use the smaller table as the right child, and the estimated cardinality remained unchanged after fixing the implementation of the operation.

Case study. Listing 3 shows issue #56714 in MySQL as an illustrative example of a performance issue caused by cardinality estimation. According to the issue report, this performance issue incurs a slowdown of execution time from 0.01 seconds to 3.02 seconds. Column b in table test uses an index, but the query in line 13

does not correctly use the index, incurring a **FULL TABLE SCAN**, which is inefficient. Although the root cause for this performance is in index selection, not in the cardinality estimator, the suboptimal index selection affects the estimated cardinality as the **FULL TABLE SCAN** is expected to scan more rows than the **INDEX SCAN**.

> Performance issues can arise from inefficient operations, flawed cardinality estimators, and inefficient query plans. The latter two causes can be found by unexpected estimated cardinalities.

4 APPROACH

We propose *CERT*, a novel technique for testing cardinality estimation. The core idea is that a given query should not have a lower estimated cardinality than a more restrictive query derived from it. We term this property *cardinality restriction monotonicity* and expect that DBMSs adhere to it in practice. *CERT* is a simple black-box technique, making it widely applicable in practice.

Method overview. Figure 1 shows an overview of *CERT* based on the running example from Listing 1. Given a randomly generated query in step ①, we derive another more restrictive query in step ② and retrieve both queries' query plans. Then, if both query plans are structurally similar in step ③, we validate the *cardinality restriction monotonicity* property in step ④; we expect the less restrictive query from step ① to return at least as high estimated cardinality as the more restrictive query from step ②. Any discrepancy is considered a performance issue. We perform the structural similarity checking in step ③ based on the observation that a more restrictive query can result in a significantly different query plan, whose estimated cardinalities are not comparable. Next, we give a detailed explanation of each step.

4.1 Database and Query Generation

We require a database state and a query for testing. Both database state and query can be manually provided or generated. Common generation-based methods include mutation-based methods [30, 67] and rule-based generation methods [42–44, 53]. How to generate database states and queries is not a contribution of this paper, and our approach can be paired with any database state and query generation method. For example, in Listing 1, we could randomly generate a database state in lines 1–4 and a query in line 8.

Before executing queries on the generated database state, we execute **ANALYZE** statements on each table to guarantee that the data statistics are up to date. For Listing 1, these statements are executed in lines 5–6.

4.2 Query Restriction

Given a query, we derive a more restrictive query based on two insights. First, for the clauses that we considered, adding a clause to a query makes the query more restrictive except for the **JOIN** clause. Second, given an already existing clause, we can modify the clause or its predicate to obtain a more restrictive query. We considered the SQL features shown in Listing 2 and propose at least one rule for each feature, yielding the 12 rules shown in Table 2. Since the **JOIN** clause, which specifies two tables or views to be joined, is a major factor influencing the queries' run time [26], 5 of the 12 rules relate
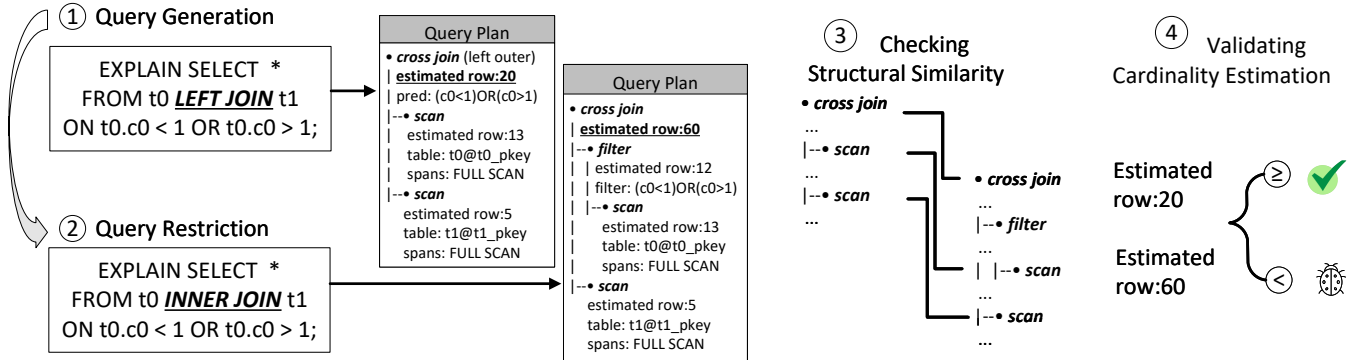
Figure 1: Overview of *CERT*.

to them. Our rules are not exhaustive; we believe that practitioners could propose additional rules depending on their testing focus.

Rule overview. In Table 2, *CERT* derives a *Target* statement from a *Source* statement by applying a restriction on the shown *Clause*, as demonstrated through *Example*. <Predicate> refers to a boolean expression and <Natural number> to a natural number. These examples are based on a database state with two tables t0 and t1, both of which have only one column c0. For each test to be generated, we randomly choose a SQL clause, of which one or more rules are randomly applied to restrict a query. In Figure 1, we choose the *JOIN* clause and apply only rule 1, which replaces a **LEFT JOIN** with an **INNER JOIN** in the **JOIN** clause of a query.

JOIN clause. Our key insight for testing the **JOIN** clause is the partial inequality relationship in terms of cardinalities between different kinds of joins. For a fixed join predicate, the following inequalities for the different joins' cardinalities hold: **INNER JOIN** ≤ **LEFT JOIN/RIGHT JOIN** ≤ **FULL JOIN** ≤ **CROSS JOIN**. Figure 2 illustrates this using a **JOIN** diagram [11] on two tables, each of which has three rows, with the same color denoting the rows that can be matched in the **JOIN** predicate. As determined by the SQL standard, **INNER JOIN** fetches rows that have matching values in both tables; **LEFT JOIN/RIGHT JOIN** fetch all rows from the left/right table and the matching rows from the respectively other table; **FULL JOIN** fetches all rows from both tables; **CROSS JOIN** fetches all possible combinations of all rows from both tables without an **ON** clause. A corner case for rule 5 concerning the **CROSS JOIN** is that this join may fetch fewer rows than **FULL JOIN** if either of the tables is empty, in which case **CROSS JOIN** fetches zero rows. To avoid potential false alarms, we ensure that each table contains at least one row.

WHERE clause. For the **WHERE** clause, our insight is that we can restrict the predicate that is used for filtering rows. If the query contains no **WHERE** clause, we restrict the query by adding one with a random predicate. If the **WHERE** clause's predicate has an **OR** operator, we restrict it by removing either of the **OR**'s operands. Otherwise, we add an **AND** operator with a randomly generated predicate. Restricting predicates would also apply to testing **JOIN** clauses; in this work, we aimed to introduce the general idea behind *CERT* and illustrate it on a small set of promising rules. We believe that practitioners who adopt the approach will propose many additional rules.
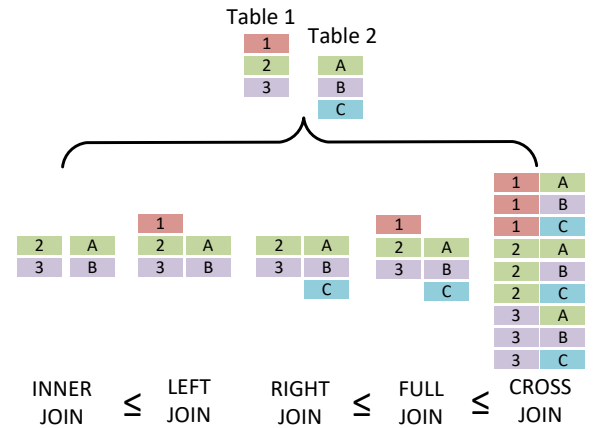

Figure 2: The inequality relationships of estimated cardinalities in the JOIN clause with an example to join two tables.

Other SQL clauses. A query can be restricted by a **DISTINCT** clause, which should fetch no more rows than the same query without such a clause, or by replacing its **ALL** clause. Similarly, a query without **GROUP BY** or **HAVING** can be restricted by adding such clauses along with any predicate. A **LIMIT** clause can be added, or a lower limit can be replaced with a higher limit.

4.3 Checking for Structural Similarity

Even for similar queries, DBMSs may create significantly different query plans. In such cases, the estimated cardinalities might be calculated in different ways, and thus result in false alarms. Listing 4 shows an example of this problem. The only difference between the two queries in lines 8–9 is that **FULL JOIN** is used in the first query and **RIGHT JOIN** is used in the second query. Their estimated cardinalities are 2 and 3 respectively. Based on rule 4 alone, this discrepancy would constitute a performance issue; however, consider the query plans in Listing 4. In lines 11–18, the left part is the query plan of the first query, and the right part is that of the second query. For the first query with **FULL JOIN**, the sequence of operations is **filter**, **cross join**, **scan**, **scan** in which the operation **filter** is applied *after* the operation **cross join**. For the second query with

Table 2: The rules to restrict queries.

	Clause	Source	Target	Example
1	JOIN	LEFT JOIN	INNER JOIN	SELECT * FROM t0 ~~LEFT~~ INNER JOIN t1 ON ...;
2	JOIN	RIGHT JOIN	INNER JOIN	SELECT * FROM t0 ~~RIGHT~~ INNER JOIN t1 ON ...;
3	JOIN	FULL JOIN	LEFT JOIN	SELECT * FROM t0 ~~FULL~~ LEFT JOIN t1 ON ...;
4	JOIN	FULL JOIN	RIGHT JOIN	SELECT * FROM t0 ~~FULL~~ RIGHT JOIN t1 ON ...;
5[†]	JOIN	CROSS JOIN	FULL JOIN	SELECT * FROM t0 ~~CROSS~~ FULL JOIN t1;
6	SELECT	ALL	DISTINCT	SELECT ~~ALL~~ DISTINCT * FROM t0;
7	GROUP BY	<Empty>	<Predicate>	SELECT * FROM t0 GROUP BY c0 ;
8	HAVING	<Empty>	<Predicate>	SELECT * FROM t0 GROUP BY c0 HAVING c0>0 ;
9	WHERE	<Empty>	<Predicate>	SELECT * FROM t0 WHERE c0>0 ;
10	WHERE	<Predicate>	<Predicate> AND <Predicate>	SELECT * FROM t0 WHERE c0>0 AND c0!=8 ;
11	WHERE	<Predicate> OR <Predicate>	<Predicate>	SELECT * FROM t0 WHERE c0>0 ~~OR c0!=8~~ ;
12	LIMIT	<Natural number>	<Natural number> - <Natural number>	SELECT * FROM t0 LIMIT ~~10~~ 5 ;

[†] Rule 5 holds when both tables are not empty.

Listing 4: An example of query plans that are not structurally similar, which is why we exclude them for testing.

```
1   CREATE TABLE t0 (c0 INT);
2   CREATE TABLE t1 (c0 INT, c1 INT);
3   INSERT INTO t1 VALUES(1,2), (3,4), (5,6), (NULL, NULL);
4   INSERT INTO t0 VALUES(1), (2);
5   ANALYZE t0;
6   ANALYZE t1;
7
8   EXPLAIN SELECT * FROM t0 FULL JOIN t1 ON t1.c1 IN (t1.c1)
        WHERE CASE WHEN t1.rowid > 2 THEN false ELSE
        t1.c1=1 END; -- estimated rows: 2
9   EXPLAIN SELECT * FROM t0 RIGHT JOIN t1 ON t1.c1 IN (t1.c1)
        WHERE CASE WHEN t1.rowid > 2 THEN false ELSE
        t1.c1=1 END; -- estimated rows: 3
10  -----------------------------------------------
11  • filter                  • cross join(right)
12  | estimated row:2         | estimated row:3
13  |-• cross join(full)      |-• scan (t0)
14    | estimated row:6       |   estimated row:2
15    |-• scan (t1)           |-• filter
16    |   estimated row:4     | | estimated row:1
17    |-• scan (t0)           |-• scan (t1)
18        estimated row:2         estimated row:4
```

RIGHT JOIN, the sequence of operations is `cross join`, `scan`, `filter`, `scan`, in which the operation `filter` is applied *before* the operation `cross join`. The difference is due to a SQL optimization mechanism called *predicate pushdown* [28], which moves a filter to be executed before joining two tables and is applied in the second query. The predicate pushdown does not affect the final result, but can reduce the estimated cardinalities to be joined in the operation **RIGHT JOIN**, which is more efficient. However, because the predicate is pushed down, the structure of the query plans changes. Therefore, the estimated cardinalities are calculated in a different manner than that of the first query, and the developers consider the estimated cardinalities of both query plans as incomparable.[7] In Listing 4, the estimated cardinalities of both operations **FULL JOIN** and **RIGHT JOIN** are calculated as the sum of the estimated cardinalities in the last

---

[7]https://github.com/cockroachdb/cockroach/issues/89060

step and the operation `filter` is calculated as one-third of the estimated cardinalities in the last step. The estimated cardinality of the first query plan is calculated by $(2 + 4)/3 = 2$, while that of the second query plan is calculated by $4/3 + 2 = 3$.

We identify comparable estimated cardinalities by checking for *structural similarity*. In a query plan P, the operation of a node N is denoted as O_N. Suppose N has k children, denoted as $C_1, C_2, \ldots, C_k$. The node's flattened operation sequence is an array obtained by concatenating the flattened child nodes: $\text{flatten}(N) = [O_N, \text{flatten}(C_1), \text{flatten}(C_2) \ldots \text{flatten}(C_k)]$. The flattening of the root node is also denoted as $\text{flatten}(P)$. For a pair of query plans P_a and P_b, we define both are *structurally similar* only if $ED(\text{flatten}(P_a), \text{flatten}(P_b)) <= 1$, in which ED represents the edit distance [34], a common way of quantifying the dissimilarity of two strings, of both operation sequences. For example, in the query plans of Figure 1, the sequences of operations are [`cross join`, `scan`, `scan`] and [`cross join`, `filter`, `scan`, `scan`]. The first sequence can be edited to the second sequence by inserting a `filter` only, and *vice versa*. Both query plans are structurally similar and we validate the *cardinality restriction monotonicity* property. If they are not structurally similar, we continue testing with a new query. The calculation is based on sequences rather than trees, because the computation of the edit distance of trees was shown to be NP-hard [50].

4.4 Validating Cardinality Estimation

Finally, we validate the *cardinality restriction monotonicity* property on the estimated cardinalities extracted from the query plans. If the estimated cardinality of the original query is lower than that of the more restrictive query, we report the query pair as an issue. In Figure 1, the estimated cardinality of the original query is 20, which is lower than that of the other restricted query, which is 60. This indicates an unexpected result, and we report both queries to developers. Recall that we deem the estimated cardinality in the root operation of the query plan as the estimated cardinality of the query and ignore the estimated cardinalities of other operations.

5 EVALUATION

To evaluate the effectiveness and efficiency of *CERT* in finding performance issues through estimated cardinalities, we implemented *CERT* in *SQLancer*,[8] which is an automated testing tool for DBMSs, and, based on our prototype *SQLancer+CERT*, we sought to answer the following questions:

Q.1 Effectiveness. Can *CERT* identify previously unknown issues?

Q.2 Historic Bugs. Can *CERT* identify historic performance issues?

Q.3 Efficiency. How does CERT compare to the state-of-the-art approach in terms of accuracy and efficiency?

Q.4 Sensitivity. Which rules proposed in Table 2 contribute to finding issues? Do DBMSs adhere to the *cardinality restriction monotonicity* property in practice?

Implementation. We reused *SQLancer* and its implementation of a rule-based random generation method to generate queries and database states. For each table, we generated 100 `INSERT` statements and ensured that every table contained at least one row. Then, the generated queries and database states are passed to *CERT* for validating the *cardinality restriction monotonicity* property. The core logic of *CERT* is implemented in only around 200 lines of Java code for each DBMS, suggesting that its low implementation effort might make the approach widely applicable. We used pattern matching to extract operations and estimated cardinalities, and implemented the *Dynamic Programming* (DP) algorithm [54] to calculate the structural similarity.

Tested DBMSs. We tested the same DBMSs, MySQL, TiDB, and CockroachDB as we studied in Section 3. For Q1, Q3, and Q4, we used the latest available development versions (MySQL: 8.0.31, TiDB: 6.4.0, CockroachDB: 22.2.0). For Q3, in an attempt of a fairer comparison to *AMOEBA*, we chose the historical version of CockroachDB 20.2.19, which is the version that *AMOEBA* [31] tested.

Baselines. To the best of our knowledge, no existing work can be applied to specifically test cardinality estimation. The most closely related work is *AMOEBA*, which finds performance issues in query optimizers. We did not consider *APOLLO* [21], because it finds only performance regressions. Ensuring a fair comparison with *AMOEBA* is challenging, as the approaches are not directly comparable. *AMOEBA* validates that semantic-equivalent queries exhibit similar performance characteristics, while *CERT* validates the *cardinality restriction monotonicity* property. Furthermore, both tools support a different set of DBMSs; *AMOEBA* supports CockroachDB and PostgreSQL, while *CERT* supports CockroachDB, TiDB, and MySQL. Thus, we performed the comparison in Q3 using only CockroachDB, which is supported by both tools.

Experimental infrastructure. We conducted all experiments on a desktop computer with an Intel(R) Core(TM) i7-9700 processor that has 8 physical cores clocked at 3.00GHz. Our test machine uses Ubuntu 20.04 with 8 GB of RAM, and a maximum utilization of 8 cores. We run all experiments 10 runs for statistical significance.

Q.1 Effectiveness

Method. We ran *SQLancer+CERT* to find performance issues. Each automatically generated issue report usually includes many SQL statements, making it challenging for developers to analyze the root reason for the issue. To alleviate this problem and better demonstrate the underlying reasons for these issues, we adopted delta debugging [65] to minimize test cases before reporting them to developers. The steps to minimize the test case are 1) incrementally removing some of the SQL statements in the test case and 2) ensuring that the *cardinality restriction monotonicity* property is still violated. After submitting the issue reports with minimized test cases to developers, we submitted follow-up issues only if we believed them to be unique, such as those identified by different rules than previous issues, to avoid duplicate issues. MySQL has its own issue-tracking system, and developers add a label *Verified* for the issues that they have confirmed. TiDB and CockroachDB use GitHub's issue tracker. TiDB's developers assign labels, such as affected versions and modules; we considered the issue as *Confirmed* after such a label was assigned. CockroachDB's developers typically directly replied whether they planned on fixing the issue which we consider *Fixed* or whether the issue was considered a false alarm. In some cases, they added a *Backlogged* label to indicate that they would investigate this issue in the future. For all DBMSs, based on historic reports, we observed that, typically, developers directly reject duplicate issue reports.

Results. Table 3 shows the unique issues that *CERT* found in three tested DBMSs. The *Bug ID* column shows the bug id in respective bug trackers. The *Version* column shows the versions or git commits of the DBMSs in which we found corresponding issues. The *Rules* column shows which rules identified this issue. In total, we have reported 13 unique issues on unexpected estimated cardinalities. 9 issues have been confirmed by developers in three days, 2 issues have been fixed in one week, and 2 issues were backlogged. No false alarm was generated. We speculate that many confirmed bugs remain unfixed, because 1) fixing performance issues requires comprehensive consideration which usually consumes much time, and 2) performance issues might have lower priority than other issues, such as correctness bugs, which cause a query to compute an incorrect result. Among all 13 unique issues, the only known issue that we found in CockroachDB had been backlogged for around 10 months since it was first found, and our test case clearly demonstrated the root reason for the issue, which allowed developers to quickly fix it. Apart from the reported issues, *CERT* continuously generates more than ten issue reports per minute. We did not report the additional bug-inducing test cases to the developers to avoid burdening them, because deciding their uniqueness would be challenging. Therefore, we believe that *CERT* could help identify additional performance issues in the future. Overall, all issues were exposed in various SQL clauses and predicates, which may imply no common issues across tested DBMSs. We give two examples of minimized test cases to explain the issues we found as follows.

Issue #108852 identified by rule 6. Listing 5 shows a test case exposing issue #108852 in MySQL. Rule 6, which replaces `ALL` with `DISTINCT` in Table 2, exposed this issue. In Listing 5, the first query

Table 3: The issues found by *CERT*. The red strike-through text is removed, while the content in green color is added.

DBMS	Bug ID	Version	Rules	Modifications to the query	Status
MySQL	108833	8.0.31	9	... WHERE t0.c0 > t0.c1 ...	Verified
MySQL	108851	8.0.31	9	... WHERE t1.c1 BETWEEN (SELECT 1 WHERE FALSE) AND (t1.c0) ...	Verified
MySQL	108852	8.0.31	6	... DISTINCT ...	Verified
TiDB	38319	51a6684f	11	...WHERE (TRUE) ~~OR(TO_BASE64(t0.c0))~~ ...	Confirmed
TiDB	38747	3ef8352a	7	... GROUP BY t0.c0 ...	Confirmed
TiDB	38479	3ef8352a	3 & 5	... ~~CROSS~~ LEFT JOIN...	Confirmed
TiDB	38482	3ef8352a	8	... HAVING (t1.c0)REGEXP(NULL) ...	Confirmed
TiDB	38665	6c55faf0	2	... ~~RIGHT~~ INNER JOIN...	Confirmed
TiDB	38721	6c55faf0	9	... WHERE v0.c2 ...	Confirmed
CockroachDB	88455	7cde315d	1	... ~~LEFT~~ INNER JOIN...	Fixed
CockroachDB	89161	f188d21d	11	...WHERE (t0.c0 IS NOT NULL) ~~OR (1 < ALL (t0.c0 & t0.c0))~~ ...	Fixed (Known)
CockroachDB	89462	81586f62	8	... HAVING (t1.c0 ::CHAR) = 'a' ...	Backlogged
CockroachDB	90113	fbfb71b9	2	... ~~RIGHT~~ INNER JOIN...	Backlogged

Listing 5: Rule 6, which identified #108852 in MySQL, replaces ALL with DISTINCT.

```
1  CREATE TABLE t0(c0 INT, c1 INT UNIQUE) ;
2  INSERT INTO t0 VALUES(-1, NULL),(1, 2),(NULL, NULL),(3,
       4);
3  ANALYZE TABLE t0 UPDATE HISTOGRAM ON c0, c1;
4
5  EXPLAIN SELECT ALL t0.c0 FROM t0 WHERE t0.c1; --
       estimated rows: 3
6  EXPLAIN SELECT DISTINCT t0.c0 FROM t0 WHERE t0.c1; --
       estimated rows: 4
```

Listing 6: Rule 11, which identified #89161 in CockroachDB, removes either operand of an OR expression.

```
1  CREATE TABLE t0 (c0 INT);
2  INSERT INTO t0 VALUES (1), (2), (3), (4), (5), (6), (7),
       (8), (9), (10);
3  ANALYZE t0;
4
5  EXPLAIN SELECT t0.c0 FROM t0 WHERE
       (t0.c0 IS NOT NULL) OR (1 < ALL (t0.c0, t0.c0)); --
       estimated rows: 3
6  EXPLAIN SELECT t0.c0 FROM t0 WHERE (t0.c0 IS NOT NULL); --
       estimated rows: 10
```

in line 5 fetches the rows including duplicate rows, while the second query in line 6 excludes duplicate rows, so the cardinality of the second query should be no more than that of the first query. However, the estimated cardinality of the second query is greater than that of the first query, which is unexpected. Suppose a query q with **ALL** is a subquery of another query Q with **DISTINCT**, this issue affects whether **DISTINCT** should be pushed down to the execution of q for an efficient query plan that aims to retrieve the fewest rows from q. This issue was confirmed by the MySQL developers already three hours after we reported it.

Listing 7: The performance improvement by fixing our found issues #88455 and #89161.

```
1  CREATE TABLE t0 (c0 INT);
2  CREATE TABLE t1 (c0 INT);
3  CREATE TABLE t2 (c0 INT);
4  INSERT INTO t0 SELECT * FROM generate_series(1,1000);
5  INSERT INTO t1 SELECT * FROM generate_series(1001,2000);
6  INSERT INTO t2 SELECT * FROM generate_series(1,333100);
7  ANALYZE t0;
8  ANALYZE t1;
9  ANALYZE t2;
10
11 SELECT COUNT(*) FROM t0 LEFT OUTER JOIN t1 ON t0.c0<1 OR
       t0.c0>1 FULL JOIN t2 ON t0.c0=t2.c0; -- 399ms →
       321ms
12 SELECT COUNT(*) FROM t0 LEFT JOIN t1 ON t0.c0>0 WHERE
       (t0.c0 IS NOT NULL) OR (1 < ALL(t0.c0, t0.c0)); --
       131ms → 109ms
```

Issue #89161 identified by rule 11. Listing 6 shows a test case exposing issue #89161 in CockroachDB by rule 11, which removes either operand of an OR expression. The predicate (t0.c0 IS **NOT NULL**) in the **WHERE** clause of the second query should fetch no more rows than the predicate (t0.c0 IS **NOT NULL**)**OR** (1 < **ALL** (t0.c0, t0.c0)) of the first query. However, the estimated cardinality of the second query is greater than that of the first query, which is unexpected. This issue was caused by a buggy logic to handle the **OR** clause. In CockroachDB, given predicates A and B, the estimated cardinality of predicate A **OR** B is calculated by: $P(A \ OR \ B) = P(A) + P(B) - P(A \ AND \ B)$. However, when A and B depend on the same table or column, the estimated cardinality is unexpected. We found this issue by rule 11 in Table 2. Although this issue was known, it had been backlogged for around 10 months since it was first found. When we reported our test case, the developer opened a pull request in their git repository to fix it after three days.

Listing 8: Issue #56714 violates the *cardinality restriction monotonicity*.

```
1  ...
2  EXPLAIN SELECT MAX(a) FROM test; -- estimated rows: 1
3  EXPLAIN SELECT MAX(a) FROM test WHERE b=0; -- estimated
       rows: 500190
```

Performance analysis. To investigate the extent to which the issues we found affect performance, we evaluated the query performance of the fixed issues #88455 and #89161 on a test case as shown in Listing 7 that involves joining multiple tables. We could not consider unfixed issues, as it would be unclear how to determine the potential speedup. We executed both queries in lines 12 and 13 before and after the fixes of #88455 and #89161 respectively. After executing either query ten times, we found that the fixes improved the performance by an average of 20% and 17%, respectively. This improvement is due to the more accurate estimated cardinality which allows for more optimal joining orders.

> Using *CERT*, we have found 13 unique issues in MySQL, TiDB, and CockroachDB. The fixes improve query performance by 19% on average.

Q.2 Historic Bugs

Method. To evaluate whether *cardinality restriction monotonicity* is sufficiently general to identify previous performance issues that we identified in Table 1, we attempted using *CERT* to identify all three performance issues whose fixes changed the estimated cardinalities, namely issues #61631, #56714, and #9067. Specifically, based on the queries in the issue reports, we followed step ② in Figure 1 to randomly construct 10,000 pairs of queries. Then, we checked whether any pair violated the *cardinality restriction monotonicity* before the fix, and adhered to the *cardinality restriction monotonicity* after the fix. If so and both query plans are structurally similar, we concluded that *cardinality restriction monotonicity* could have identified the performance issue.

Results. All three previous performance issues caused by cardinality estimation can be found by *CERT*. For example, considering the performance issue #56714 in Listing 3, Listing 8 shows the pair of queries that *CERT* produces to identify the performance issue. The second query has an additional `WHERE` clause compared to the first query, so the estimated cardinality of the second query should be no more than that of the first query. However, due to incorrect usage of the index in column `b`, the second query scans all rows and has a higher estimated cardinality. After the fix, the estimated cardinality of the second query decreases to 1.

> The *cardinality restriction monotonicity* can identify historical performance issues caused by cardinality estimation.

Q.3 Efficiency

Accuracy. We evaluated whether *CERT* has higher accuracy in confirmed issues than *AMOEBA*. We evaluated this aspect based on the observation that around five in six reported bugs by *AMOEBA*

Table 4: The number of all (All) and confirmed or fixed (C/F) unique performance issues.

DBMS	CERT			AMOEBA		
	All	C/F	%	All	C/F	%
MySQL	3	3	100%	-	-	-
TiDB	6	6	100%	-	-	-
CockroachDB	4	2	50%	25	6	24%
Sum:	13	11	85%	25	6	24%

were false alarms. A high rate of false alarms significantly limits the applicability of an automated testing technique. Recall that it is challenging to make a fair comparison as *CERT* and *AMOEBA* find different kinds of issues affecting performance.

Results. Table 4 shows the number of all and confirmed/fixed unique performance issues found by *CERT* and *AMOEBA*. The authors of *AMOEBA* reported 25 issues in CockroachDB, but only 6 issues (24% accuracy) were confirmed or fixed by developers. In comparison, for *CERT*, 50% of issues in CockroachDB and 100% of issues in MySQL and TiDB were confirmed or fixed. For CockroachDB, *CERT* found fewer performance issues than *AMOEBA*, because we did not report all found issues to avoid duplicate reports. Overall, these results suggest that *CERT* has high accuracy and is a practical technique for finding relevant performance issues. Despite these promising results, on a conceptual level, similar to *AMOEBA*, we cannot ensure that the performance issues would be considered as such by the developers.

Throughput. We evaluated whether *CERT* has a higher testing throughput than *AMOEBA*. State-of-the-art benchmarks and approaches, such as TPC-H [33], *AMOEBA* [31], and *APOLLO* [21] execute queries, which results in relatively low throughput. Therefore, we expect that *CERT* will validate more queries per second. To evaluate this, we determined the average number of test cases per second processed by *SQLancer+CERT* and *AMOEBA* in one hour.

Results. In CockroachDB, on average across 10 runs and one hour, *CERT* validates 714.54 test cases while *AMOEBA* exercises 1.85 test cases per second. This suggests a 386× performance improvement over *AMOEBA*. Recall that the throughput results are not directly comparable, as different approaches can find different kinds of issues. In addition, *CERT* is immune to performance fluctuation, because the estimated cardinalities are not affected by execution time. Therefore, *CERT* yields the same results in different hardware and network environments.

> *SQLancer+CERT* validates 386× more test cases than *AMOEBA* across one hour and 10 runs on average, and is more than twice as accurate as *AMOEBA*.

Q.4 Sensitivity

Sensitivity of rules. We evaluated which rules presented in Table 2 contribute to finding the issues in Table 3. Specifically, we recorded which rules were applied to the bug-inducing test cases that we reported. We considered reported bug-inducing test cases,

Table 5: The average number of all queries (Queries), the queries that violate *cardinality restriction monotonicity* (Violations), and the geometric mean of percentage (%) of queries that violate the property across 10 runs and one hour.

DBMS	Queries (#)	Violations (#)	Violations (%)
MySQL	6,371,222	30,841	0.28%
TiDB	2,895,203	8,108	0.27%
CockroachDB	1,306,807	661	0.05%
		Average:	0.2%

rather than all test cases—recall that *SQLancer+CERT* still reports violations when being run—as we expect the reported issues to be unique based on the developers' verdicts. Table 3 shows the rules applied to the corresponding bug-inducing test cases. Overall, 9 out of 12 rules have found at least one issue. Rule 9, which adds a predicate to `WHERE` clause, found the most issues, namely three. We believe that this is because the predicates in `WHERE` clause can vary significantly and thus be diverse and have a higher possibility of exposing issues. No issue was found by rules 4, 10, and 12. Rule 4 applies to the `JOIN` clause in which other rules found several issues. Similarly, we believe that rule 10, which restricts a `WHERE` clause by an `AND` operator would find issues after the issues found by other rules applied to `WHERE` clause are fixed. Rule 12 applies to the `LIMIT` clause, which simply returns up to as many rows as specified in its argument. We explain that the simplicity of `LIMIT` explains that we have found no issues in its handling.

Sensitivity of cardinality restriction monotonicity. We expect that any violation of the *cardinality restriction monotonicity* property indicates a potential issue. To more thoroughly assess our hypothesis, we examined how many queries among all tested queries violate the *cardinality restriction monotonicity* property. If only a small portion of queries violate it, DBMSs are likely to adhere to the *cardinality restriction monotonicity* property, and any violation warrants further investigation. Otherwise, the property may be not meaningful for developers. Table 5 shows the average number of all queries, the queries that violate *cardinality restriction monotonicity* property, and the geometric mean of the percentage of queries that violate the property across 10 runs and one hour. Overall, 0.2% of all generated queries violate the *cardinality restriction monotonicity* property. All three DBMSs, MySQL, TiDB, and CockroachDB, exhibit a similar rate of *cardinality restriction monotonicity* violations. The results demonstrate that DBMSs typically comply with the *cardinality restriction monotonicity* property, as more than 99.5% of generated queries do not violate it.

6 DISCUSSION

We discuss some key considerations on the design of *CERT*, its characteristics, as well as the evaluation's results.

Evaluating performance gains. It is challenging to measure the overall performance gain that fixing the issues reported by *CERT* could achieve in practice. One issue is that measuring the overall performance impact might be misleading, because many other components in query optimizers can affect performance as well. For example, research on the Join Order Benchmark [26] demonstrated that a worse cardinality estimator might lead to better performance due to the issues in other components. In addition, 9 issues were confirmed, but remained unfixed. Since we lack domain knowledge to address the underlying issues, we cannot determine the performance gains that fixing these issues would cause.

Generality. While we focused on important SQL clauses in this work, the 12 rules shown in Table 2 are not comprehensive. They could be extended to cover various other features, such as window functions. Additionally, the *cardinality restriction monotonicity* property might be applicable also to other kinds of data models than the traditional relational data model. For example, Neo4J, a graph DBMS, also uses a concept similar to query plans—termed execution plans—and cardinality estimation[9] (*EstimatedRows* field in execution plans). Its query optimization also depends on cardinality estimation, which we expect to comply with the *cardinality restriction monotonicity* property. More work is required to explore *cardinality restriction monotonicity* in other DBMSs in the future.

Threats to validity. Our evaluation results face potential threats to validity. One concern is internal validity, that is, the degree to which our results minimize systematic error. *CERT* validates test cases that are randomly generated by *SQLancer*. The randomness process may limit the reproducibility of our results. To mitigate the risk, we repeated all experiments 10 times to gain statistical significance. Another concern is external validity, that is, the degree to which our results can be generalized to and across other DBMSs. We selected various types of DBMSs including different purposes (community-developed: MySQL and company-backed: TiDB and CockroachDB) and languages (C/C++: MySQL and Go: TiDB and CockroachDB). These DBMSs have been widely used in prior research [30, 43, 44]. Given that DBMSs provide similar functionality and features, we are confident that our results generalize to many DBMSs. The last concern is construct validity, that is, the degree to which our evaluation accurately assesses what the results are supposed to. *CERT* found 13 unique performance issues, but only 2 issues had been fixed posing the threat that developers might not fix them in the future or might deem them less important. To address this threat, we communicated with the TiDB developers, who informed us that they plan to fix the issues and indicated an interest in using *SQLancer+CERT*.

7 RELATED WORK

Automatically testing for performance issues. The most related strand of research is on finding performance issues in DBMSs automatically. Jung *et al.* proposed *APOLLO* [21], which compares the execution times of a query on two versions of a database system to find performance regression bugs. Liu *et al.* proposed *AMOEBA* [31], which compares the execution time of a semantically-equivalent pair of queries to identify an unexpected slowdown. In contrast, *CERT* specifically tests cardinality estimation, which is most critical for query optimization [26]. Thus, we believe that issues found by *CERT* might be most relevant for DBMSs' performance. Unlike these works, *CERT* avoids executing queries by inspecting only query

---

[9]https://neo4j.com/docs/cypher-manual/current/execution-plans/#execution-plan-introduction

plans, allowing for higher throughput and making the approach robust against unanticipated performance fluctuations.

Performance benchmarking. Benchmarking DBMSs is a common practice to identify performance regressions, and to continuously improve the DBMSs' performance on a set of benchmarks. *TPC-H* [33] and *TPC-DS* [51] are the most popular benchmarks and are considered the industry standard. Boncz *et al.* studied and identified 28 "chokepoints" (*i.e.*, optimization challenges) of the TPC-H benchmark [4]. Poess *et al.* modified and analyzed the TPC-DS benchmark [38] to measure SQL-based big data systems. Karimov *et al.* proposed a benchmark for stream-data-processing systems [22]. Boncz *et al.* proposed an improved TPC-H benchmark,*JCC-H* [3], which introduces *Join-Crossing-Correlations* (JCC) to evaluate the scenarios where data in one table can affect the behaviors of operations involving data in other tables. Leis *et al.* proposed the *Join Order Benchmark (JOB)* [26], which uses more complex join orders. Raasveldt *et al.* described common pitfalls when benchmarking DBMSs and demonstrated how they can affect a DBMS's performance [40]. *CERT* is complementary to benchmarking; while benchmarking focuses on workloads deemed relevant for users, *CERT* can find performance issues through the lens of cardinality estimation even on previously unseen workloads.

DBMS testing. Besides testing DBMSs' performance, automated testing approaches have been proposed to find other types of bugs. Many fuzzing works on finding security-relevant bugs, such as memory errors, were proposed. SQLSmith [53], Griffin [14], Dyn-SQL [20], and ADUSA [31] used grammar-based methods to generate test cases for detecting memory errors. Squirrel [67], inspired by grey-box fuzzers such as AFL [52], used code coverage as guidance to find memory errors. Various approaches have been proposed to find logic bugs in DBMSs. The PQS [44], NoREC [42], and TLP [43] oracles detect logic bugs in the implementation of `SELECT` statements. DQE [47] detects logic bugs in `UPDATE` and `INSERT` statements. Transactional properties have also been tested. Cobra [49] used formal methods to verify the serializability of executions in key-value stores. Troc [9] defines an oracle to detect isolation bugs in transactional procedures. Similar to testing approaches that find logic bugs and test properties, our core contribution is a new test oracle, which, however, finds performance issues that violate the *cardinality restriction monotonicity* property.

Cardinality estimation. Various cardinality-estimation techniques have been proposed. Han *et al.* [17] comprehensively evaluated various algorithms for cardinality estimation, of which we describe important ones below. PostgreSQL [55] and MultiHist [39] applied one-dimensional and multi-dimensional histograms to estimate cardinality. Similarly, UniSample [27, 66] and WJSample [29] used sampling-based methods to estimate cardinality. Apart from these traditional approaches, machine learning-based methods have gained attention recently. MSCN [23, 64], LW-XGB [10], and UAE-Q [59] used deep neural networks, classic lightweight regression models, and deep auto-regression models to learn to map each query directly to its estimated cardinality. In addition, NeuroCard [63], BayesCard [58], DeepDB [19], and FSPN [60, 68] utilized deep auto-regression models and three probabilistic graphical models BN,

SPN, and FSPN to predict the data distribution for cardinality estimation. *CERT* is a black-box technique that could, in principle, be applied to any cardinality estimator. However, we believe that finding relevant issues in learning-based estimators that can be fixed might be challenging.

Metamorphic testing. At a conceptual level, *CERT* can be classified as a metamorphic testing approach [7, 8], which is a method to generate both test cases and validate results. Metamorphic testing uses an input I to a system and its output O to derive a new input I' (and output O'), for which a test oracle can be provided that checks whether a so-called *Metamorphic Relation* holds between O and O'. For *CERT*, I corresponds to the original query, O is the estimated cardinality, I' is the restricted query, and O' is its estimated cardinality. The metamorphic relation that we validate is the *cardinality restriction monotonicity* property. Metamorphic testing has been applied successfully in various domains [8, 45]. TLP [43] and NoREC [42] are other metamorphic testing approaches that were proposed for testing DBMSs, as discussed above.

8 CONCLUSION

We have presented *CERT*, a novel technique for finding performance issues through the lens of cardinality estimation in DBMSs. Our key idea is to, given a query, derive a more restrictive query, and validate that the restriction is reflected by the DBMSs' estimated cardinalities; we refer to this property as *cardinality restriction monotonicity*. Specifically, we validate that the original query's cardinality is at least as high as the more restrictive query. Our evaluation has demonstrated that this technique is effective. Of the 13 unique issues that we reported, 2 issues were fixed, 9 issues were confirmed, and 2 issues require further investigation. The fixes improved query performance by 19% on average. Unlike other testing approaches for performance issues, *CERT* avoids executing queries, achieving a speedup of 386× compared to the state of the art. Finally, it is readily applicable. DBMSs expose estimated cardinalities as part of query plans to users; thus, *CERT* is a black-box technique that is applicable without modifications, even if the DBMSs' source code is inaccessible to the testers. Furthermore, since no queries are executed, *CERT* is resistant to performance fluctuations. Overall, we believe that *CERT* is a useful technique for DBMS developers and testers and hope that the technique will be widely adopted in practice.

9 DATA AVAILABILITY

Our implementation and experimental data are publicly available at https://zenodo.org/records/10476847.

REFERENCES

[1] Mert Akdere, Ugur Çetintemel, Matteo Riondato, Eli Upfal, and Stanley B. Zdonik. 2012. Learning-based Query Performance Modeling and Prediction. In *IEEE 28th International Conference on Data Engineering (ICDE 2012), Washington, DC, USA (Arlington, Virginia), 1-5 April, 2012*, Anastasios Kementsietsidis and Marcos Antonio Vaz Salles (Eds.). IEEE Computer Society, 390–401. https://doi.org/10.1109/ICDE.2012.64

[2] Thomas Bach, Artur Andrzejak, Changyun Seo, Christian Bierstedt, Christian Lemke, Daniel Ritter, Dongwon Hwang, Erda Sheshi, Felix Schabernack, Frank Renkes, Gordon Gaumnitz, Jakob Martens, Lars Hömke, Michael Felderer, Michael Rudolf, Neetha Jambigi, Norman May, Robin Joy, Ruben Scheja, Sascha Schwedes, Sebastian Seibel, Sebastian Seifert, Stefan Haas, Stephan Kraft, Thomas Kroll, Tobias Scheuer, and Wolfgang Lehner. 2022. Testing Very Large Database Management Systems: The Case of SAP HANA. *Datenbank-Spektrum* 22, 3 (2022), 195–215. https://doi.org/10.1007/s13222-022-00426-x

[3] Peter A. Boncz, Angelos-Christos G. Anadiotis, and Steffen Kläbe. 2017. JCC-H: Adding Join Crossing Correlations with Skew to TPC-H. In *Performance Evaluation and Benchmarking for the Analytics Era - 9th TPC Technology Conference, TPCTC 2017, Munich, Germany, August 28, 2017, Revised Selected Papers (Lecture Notes in Computer Science, Vol. 10661)*, Raghunath Nambiar and Meikel Poess (Eds.). Springer, 103–119. https://doi.org/10.1007/978-3-319-72401-0_8

[4] Peter A. Boncz, Thomas Neumann, and Orri Erling. 2013. TPC-H Analyzed: Hidden Messages and Lessons Learned from an Influential Benchmark. In *Performance Characterization and Benchmarking - 5th TPC Technology Conference, TPCTC 2013, Trento, Italy, August 26, 2013, Revised Selected Papers (Lecture Notes in Computer Science, Vol. 8391)*, Raghunath Nambiar and Meikel Poess (Eds.). Springer, 61–76. https://doi.org/10.1007/978-3-319-04936-6_5

[5] Donald D. Chamberlin and Raymond F. Boyce. 1974. SEQUEL: A Structured English Query Language. In *Proceedings of 1974 ACM-SIGMOD Workshop on Data Description, Access and Control, Ann Arbor, Michigan, USA, May 1-3, 1974, 2 Volumes*, Gene Altshuler, Randall Rustin, and Bernard D. Plagman (Eds.). ACM, 249–264. https://doi.org/10.1145/800296.811515

[6] Surajit Chaudhuri. 1998. An Overview of Query Optimization in Relational Systems. In *Proceedings of the Seventeenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, June 1-3, 1998, Seattle, Washington, USA*, Alberto O. Mendelzon and Jan Paredaens (Eds.). ACM Press, 34–43. https://doi.org/10.1145/275487.275492

[7] Tsong Yueh Chen, S. C. Cheung, and Siu-Ming Yiu. 2020. Metamorphic Testing: A New Approach for Generating Next Test Cases. *CoRR* abs/2002.12543 (2020). arXiv:2002.12543 https://arxiv.org/abs/2002.12543

[8] Tsong Yueh Chen, Fei-Ching Kuo, Huai Liu, Pak-Lok Poon, Dave Towey, T. H. Tse, and Zhi Quan Zhou. 2018. Metamorphic Testing: A Review of Challenges and Opportunities. *ACM Comput. Surv.* 51, 1 (2018), 4:1–4:27. https://doi.org/10.1145/3143561

[9] Wensheng Dou, Ziyu Cui, Qianwang Dai, Jiansen Song, Dong Wang, Yu Gao, Wei Wang, Jun Wei, Lei Chen, Hanmo Wang, et al. 2023. Detecting isolation bugs via transaction oracle construction. In *Proceedings of International Conference on Software Engineering (ICSE)*.

[10] Anshuman Dutt, Chi Wang, Azade Nazi, Srikanth Kandula, Vivek R. Narasayya, and Surajit Chaudhuri. 2019. Selectivity Estimation for Range Predicates using Lightweight Models. *Proc. VLDB Endow.* 12, 9 (2019), 1044–1057. https://doi.org/10.14778/3329772.3329780

[11] Lukas Eder. 2022. Say NO to Venn Diagrams When Explaining JOINs. https://blog.jooq.org/say-no-to-venn-diagrams-when-explaining-joins/. Accessed: 2022-11-15.

[12] Pit Fender and Guido Moerkotte. 2013. Counter Strike: Generic Top-Down Join Enumeration for Hypergraphs. *Proc. VLDB Endow.* 6, 14 (2013), 1822–1833. https://doi.org/10.14778/2556549.2556565

[13] Pit Fender, Guido Moerkotte, Thomas Neumann, and Viktor Leis. 2012. Effective and Robust Pruning for Top-Down Join Enumeration Algorithms. In *IEEE 28th International Conference on Data Engineering (ICDE 2012), Washington, DC, USA (Arlington, Virginia), 1-5 April, 2012*, Anastasios Kementsietsidis and Marcos Antonio Vaz Salles (Eds.). IEEE Computer Society, 414–425. https://doi.org/10.1109/ICDE.2012.27

[14] Jingzhou Fu, Jie Liang, Zhiyong Wu, Mingzhe Wang, and Yu Jiang. 2022. Griffin: Grammar-Free DBMS Fuzzing. In *37th IEEE/ACM International Conference on Automated Software Engineering, ASE 2022, Rochester, MI, USA, October 10-14, 2022*. ACM, 49:1–49:12. https://doi.org/10.1145/3551349.3560431

[15] Jana Giceva, Gustavo Alonso, Timothy Roscoe, and Tim Harris. 2014. Deployment of Query Plans on Multicores. *Proc. VLDB Endow.* 8, 3 (2014), 233–244. https://doi.org/10.14778/2735508.2735513

[16] Goetz Graefe and Harumi A. Kuno. 2011. Modern B-tree techniques. In *Proceedings of the 27th International Conference on Data Engineering, ICDE 2011, April 11-16, 2011, Hannover, Germany*, Serge Abiteboul, Klemens Böhm, Christoph Koch, and Kian-Lee Tan (Eds.). IEEE Computer Society, 1370–1373. https://doi.org/10.1109/ICDE.2011.5767956

[17] Yuxing Han, Ziniu Wu, Peizhi Wu, Rong Zhu, Jingyi Yang, Liang Wei Tan, Kai Zeng, Gao Cong, Yanzhao Qin, Andreas Pfadler, Zhengping Qian, Jingren Zhou, Jiangneng Li, and Bin Cui. 2021. Cardinality Estimation in DBMS: A Comprehensive Benchmark Evaluation. *Proc. VLDB Endow.* 15, 4 (2021), 752–765. https://doi.org/10.14778/3503585.3503586

[18] Max Heimel, Martin Kiefer, and Volker Markl. 2015. Self-Tuning, GPU-Accelerated Kernel Density Models for Multidimensional Selectivity Estimation. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data, Melbourne, Victoria, Australia, May 31 - June 4, 2015*, Timos K. Sellis, Susan B. Davidson, and Zachary G. Ives (Eds.). ACM, 1477–1492. https://doi.org/10.1145/2723372.2749438

[19] Benjamin Hilprecht, Andreas Schmidt, Moritz Kulessa, Alejandro Molina, Kristian Kersting, and Carsten Binnig. 2020. DeepDB: Learn from Data, not from Queries! *Proc. VLDB Endow.* 13, 7 (2020), 992–1005. https://doi.org/10.14778/3384345.3384349

[20] Zu-Ming Jiang, Jia-Ju Bai, and Zhendong Su. 2023. DynSQL: Stateful Fuzzing for Database Management Systems with Complex and Valid SQL Query Generation. In *32st USENIX Security Symposium (USENIX Security 23)*.

[21] Jinho Jung, Hong Hu, Joy Arulraj, Taesoo Kim, and Woon-Hak Kang. 2019. APOLLO: Automatic Detection and Diagnosis of Performance Regressions in Database Systems. *Proc. VLDB Endow.* 13, 1 (2019), 57–70. https://doi.org/10.14778/3357377.3357382

[22] Jeyhun Karimov, Tilmann Rabl, Asterios Katsifodimos, Roman Samarev, Henri Heiskanen, and Volker Markl. 2018. Benchmarking Distributed Stream Data Processing Systems. In *34th IEEE International Conference on Data Engineering, ICDE 2018, Paris, France, April 16-19, 2018*. IEEE Computer Society, 1507–1518. https://doi.org/10.1109/ICDE.2018.00169

[23] Andreas Kipf, Thomas Kipf, Bernhard Radke, Viktor Leis, Peter A. Boncz, and Alfons Kemper. 2019. Learned Cardinalities: Estimating Correlated Joins with Deep Learning. In *9th Biennial Conference on Innovative Data Systems Research, CIDR 2019, Asilomar, CA, USA, January 13-16, 2019, Online Proceedings*. www.cidrdb.org. http://cidrdb.org/cidr2019/papers/p101-kipf-cidr19.pdf

[24] Christoph Laaber, Joel Scheuner, and Philipp Leitner. 2019. Software microbenchmarking in the cloud. How bad is it really? *Empir. Softw. Eng.* 24, 4 (2019), 2469–2508. https://doi.org/10.1007/s10664-019-09681-1

[25] Doug Laney et al. 2001. 3D data management: Controlling data volume, velocity and variety. *META group research note* 6, 70 (2001), 1.

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

[27] Viktor Leis, Bernhard Radke, Andrey Gubichev, Alfons Kemper, and Thomas Neumann. 2017. Cardinality Estimation Done Right: Index-Based Join Sampling. In *8th Biennial Conference on Innovative Data Systems Research, CIDR 2017, Chaminade, CA, USA, January 8-11, 2017, Online Proceedings*. www.cidrdb.org. http://cidrdb.org/cidr2017/papers/p9-leis-cidr17.pdf

[28] Alon Y. Levy, Inderpal Singh Mumick, and Yehoshua Sagiv. 1994. Query Optimization by Predicate Move-Around. In *VLDB'94, Proceedings of 20th International Conference on Very Large Data Bases, September 12-15, 1994, Santiago de Chile, Chile*, Jorge B. Bocca, Matthias Jarke, and Carlo Zaniolo (Eds.). Morgan Kaufmann, 96–107. http://www.vldb.org/conf/1994/P096.PDF

[29] Feifei Li, Bin Wu, Ke Yi, and Zhuoyue Zhao. 2016. Wander Join: Online Aggregation via Random Walks. In *Proceedings of the 2016 International Conference on Management of Data, SIGMOD Conference 2016, San Francisco, CA, USA, June 26 - July 01, 2016*, Fatma Özcan, Georgia Koutrika, and Sam Madden (Eds.). ACM, 615–629. https://doi.org/10.1145/2882903.2915235

[30] Yu Liang, Song Liu, and Hong Hu. 2022. Detecting Logical Bugs of DBMS with Coverage-based Guidance. In *31st USENIX Security Symposium, USENIX Security 2022, Boston, MA, USA, August 10-12, 2022*, Kevin R. B. Butler and Kurt Thomas (Eds.). USENIX Association, 4309–4326. https://www.usenix.org/conference/usenixsecurity22/presentation/liang

[31] Xinyu Liu, Qi Zhou, Joy Arulraj, and Alessandro Orso. 2022. Automatic Detection of Performance Bugs in Database Systems using Equivalent Queries. In *44th IEEE/ACM 44th International Conference on Software Engineering, ICSE 2022, Pittsburgh, PA, USA, May 25-27, 2022*. ACM, 225–236. https://doi.org/10.1145/3510003.3510093

[32] Todd Mytkowicz, Amer Diwan, Matthias Hauswirth, and Peter F. Sweeney. 2009. Producing wrong data without doing anything obviously wrong!. In *Proceedings of the 14th International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS 2009, Washington, DC, USA, March 7-11, 2009*, Mary Lou Soffa and Mary Jane Irwin (Eds.). ACM, 265–276. https://doi.org/10.1145/1508244.1508275

[33] Raghunath Othayoth Nambiar, Meikel Poess, Andrew Masland, H. Reza Taheri, Matthew Emmerton, Forrest Carman, and Michael Majdalany. 2012. TPC Benchmark Roadmap 2012. In *Selected Topics in Performance Evaluation and Benchmarking - 4th TPC Technology Conference, TPCTC 2012, Istanbul, Turkey, August 27, 2012, Revised Selected Papers (Lecture Notes in Computer Science, Vol. 7755)*, Raghunath Othayoth Nambiar and Meikel Poess (Eds.). Springer, 1–20. https://doi.org/10.1007/978-3-642-36727-4_1

[34] Gonzalo Navarro. 2001. A guided tour to approximate string matching. *ACM computing surveys (CSUR)* 33, 1 (2001), 31–88.

[35] Thomas Neumann. 2009. Query simplification: graceful degradation for join-order optimization. In *Proceedings of the ACM SIGMOD International Conference on Management of Data, SIGMOD 2009, Providence, Rhode Island, USA, June 29 - July 2, 2009*, Ugur Çetintemel, Stanley B. Zdonik, Donald Kossmann, and Nesime Tatbul (Eds.). ACM, 403–414. https://doi.org/10.1145/1559845.1559889

[36] Richard E. Pattis. 1994. Teaching EBNF first in CS 1. In *Proceedings of the 25th SIGCSE Technical Symposium on Computer Science Education, SIGCSE 1994, Phoenix, Arizona, USA, March 10-12, 1994*, Robert Beck and Don Goelman (Eds.). ACM, 300–303. https://doi.org/10.1145/191029.191155

[37] Johns Paul, Jiong He, and Bingsheng He. 2016. GPL: A GPU-based Pipelined Query Processing Engine. In *Proceedings of the 2016 International Conference on Management of Data, SIGMOD Conference 2016, San Francisco, CA, USA, June 26 - July 01, 2016*, Fatma Özcan, Georgia Koutrika, and Sam Madden (Eds.). ACM, 1935–1950. https://doi.org/10.1145/2882903.2915224

[38] Meikel Poess, Tilmann Rabl, and Hans-Arno Jacobsen. 2017. Analysis of TPC-DS: the first standard benchmark for SQL-based big data systems. In *Proceedings of the 2017 Symposium on Cloud Computing, SoCC 2017, Santa Clara, CA, USA, September 24-27, 2017*. ACM, 573–585. https://doi.org/10.1145/3127479.3128603

[39] Viswanath Poosala and Yannis E. Ioannidis. 1997. Selectivity Estimation Without the Attribute Value Independence Assumption. In *VLDB'97, Proceedings of 23rd International Conference on Very Large Data Bases, August 25-29, 1997, Athens, Greece*, Matthias Jarke, Michael J. Carey, Klaus R. Dittrich, Frederick H. Lochovsky, Pericles Loucopoulos, and Manfred A. Jeusfeld (Eds.). Morgan Kaufmann, 486–495. http://www.vldb.org/conf/1997/P486.PDF

[40] Mark Raasveldt, Pedro Holanda, Tim Gubner, and Hannes Mühleisen. 2018. Fair Benchmarking Considered Difficult: Common Pitfalls In Database Performance Testing. In *Proceedings of the 7th International Workshop on Testing Database Systems, DBTest@SIGMOD 2018, Houston, TX, USA, June 15, 2018*, Alexander Böhm and Tilmann Rabl (Eds.). ACM, 2:1–2:6. https://doi.org/10.1145/3209950.3209955

[41] Kim-Thomas Rehmann, Changyun Seo, Dongwon Hwang, Binh Than Truong, Alexander Boehm, and Dong Hun Lee. 2016. Performance Monitoring in SAP HANA's Continuous Integration Process. *SIGMETRICS Perform. Evaluation Rev.* 43, 4 (2016), 43–52. https://doi.org/10.1145/2897356.2897362

[42] Manuel Rigger and Zhendong Su. 2020. Detecting optimization bugs in database engines via non-optimizing reference engine construction. In *ESEC/FSE '20: 28th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering, Virtual Event, USA, November 8-13, 2020*, Prem Devanbu, Myra B. Cohen, and Thomas Zimmermann (Eds.). ACM, 1140–1152. https://doi.org/10.1145/3368089.3409710

[43] Manuel Rigger and Zhendong Su. 2020. Finding bugs in database systems via query partitioning. *Proc. ACM Program. Lang.* 4, OOPSLA (2020), 211:1–211:30. https://doi.org/10.1145/3428279

[44] Manuel Rigger and Zhendong Su. 2020. Testing Database Engines via Pivoted Query Synthesis. In *14th USENIX Symposium on Operating Systems Design and Implementation, OSDI 2020, Virtual Event, November 4-6, 2020*. USENIX Association, 667–682. https://www.usenix.org/conference/osdi20/presentation/rigger

[45] Sergio Segura and Zhi Quan Zhou. 2018. Metamorphic testing 20 years later: a hands-on introduction. In *Proceedings of the 40th International Conference on Software Engineering: Companion Proceeedings, ICSE 2018, Gothenburg, Sweden, May 27 - June 03, 2018*, Michel Chaudron, Ivica Crnkovic, Marsha Chechik, and Mark Harman (Eds.). ACM, 538–539. https://doi.org/10.1145/3183440.3183468

[46] Patricia G. Selinger, Morton M. Astrahan, Donald D. Chamberlin, Raymond A. Lorie, and Thomas G. Price. 1979. Access Path Selection in a Relational Database Management System. In *Proceedings of the 1979 ACM SIGMOD International Conference on Management of Data, Boston, Massachusetts, USA, May 30 - June 1*, Philip A. Bernstein (Ed.). ACM, 23–34. https://doi.org/10.1145/582095.582099

[47] Jiansen Song, Wensheng Dou, Ziyu Cui, Qianwang Dai, Wei Wang, Jun Wei, Hua Zhong, and Tao Huang. 2023. Testing database systems via differential query execution. In *Proceedings of IEEE/ACM International Conference on Software Engineering (ICSE)*.

[48] Michael Stonebraker, Sam Madden, and Pradeep Dubey. 2013. Intel "big data" science and technology center vision and execution plan. *SIGMOD Rec.* 42, 1 (2013), 44–49. https://doi.org/10.1145/2481528.2481537

[49] Cheng Tan, Changgeng Zhao, Shuai Mu, and Michael Walfish. 2020. Cobra: Making Transactional Key-Value Stores Verifiably Serializable. In *14th USENIX Symposium on Operating Systems Design and Implementation, OSDI 2020, Virtual Event, November 4-6, 2020*. USENIX Association, 63–80. https://www.usenix.org/conference/osdi20/presentation/tan

[50] Hélène Touzet. 2003. Tree edit distance with gaps. *Inf. Process. Lett.* 85, 3 (2003), 123–129. https://doi.org/10.1016/S0020-0190(02)00369-1

[51] Website. 1988. TPC-DS Benchmark. https://www.tpc.org/tpcds/. Accessed: 2022-11-15.

[52] Website. 2013. American Fuzzy Lop (AFL) Fuzzer. http://lcamtuf.coredump.cx/afl/technical_details.txt. Accessed: 2022-11-15.

[53] Website. 2015. SQLsmith. https://github.com/anse1/sqlsmith. Accessed: 2022-11-15.

[54] Website. 2020. Dynamic programming and edit distance. https://www.cs.jhu.edu/~langmea/resources/lecture_notes/dp_and_edit_dist.pdf. Accessed: 2022-11-15.

[55] Website. 2022. PostgreSQL. https://www.postgresql.org/docs/current/row-estimation-examples.html. Accessed: 2022-11-15.

[56] Website. 2022. What is Database. https://www.oracle.com/database/what-is-database. Accessed: 2022-11-15.

[57] Wentao Wu, Yun Chi, Shenghuo Zhu, Jun'ichi Tatemura, Hakan Hacigümüs, and Jeffrey F. Naughton. 2013. Predicting query execution time: Are optimizer cost models really unusable?. In *29th IEEE International Conference on Data Engineering, ICDE 2013, Brisbane, Australia, April 8-12, 2013*, Christian S. Jensen, Christopher M. Jermaine, and Xiaofang Zhou (Eds.). IEEE Computer Society, 1081–1092. https://doi.org/10.1109/ICDE.2013.6544899

[58] Ziniu Wu, Amir Shaikhha, Rong Zhu, Kai Zeng, Yuxing Han, and Jingren Zhou. 2020. BayesCard: Revitilizing Bayesian Frameworks for Cardinality Estimation. *arXiv preprint arXiv:2012.14743* (2020).

[59] Ziniu Wu, Pei Yu, Peilun Yang, Rong Zhu, Yuxing Han, Yaliang Li, Defu Lian, Kai Zeng, and Jingren Zhou. 2022. A Unified Transferable Model for ML-Enhanced DBMS. In *12th Conference on Innovative Data Systems Research, CIDR 2022, Chaminade, CA, USA, January 9-12, 2022*. www.cidrdb.org. https://www.cidrdb.org/cidr2022/papers/p6-wu.pdf

[60] Ziniu Wu, Rong Zhu, Andreas Pfadler, Yuxing Han, Jiangneng Li, Zhengping Qian, Kai Zeng, and Jingren Zhou. 2020. FSPN: A New Class of Probabilistic Graphical Model. *CoRR* abs/2011.09020 (2020). arXiv:2011.09020 https://arxiv.org/abs/2011.09020

[61] Khaled Yagoub, Peter Belknap, Benoît Dageville, Karl Dias, Shantanu Joshi, and Hailing Yu. 2008. Oracle's SQL Performance Analyzer. *IEEE Data Eng. Bull.* 31, 1 (2008), 51–58. http://sites.computer.org/debull/A08mar/yagoub.pdf

[62] Jiaqi Yan, Qiuye Jin, Shrainik Jain, Stratis D. Viglas, and Allison W. Lee. 2018. Snowtrail: Testing with Production Queries on a Cloud Database. In *Proceedings of the 7th International Workshop on Testing Database Systems, DBTest@SIGMOD 2018, Houston, TX, USA, June 15, 2018*, Alexander Böhm and Tilmann Rabl (Eds.). ACM, 4:1–4:6. https://doi.org/10.1145/3209950.3209958

[63] Zongheng Yang, Amog Kamsetty, Sifei Luan, Eric Liang, Yan Duan, Xi Chen, and Ion Stoica. 2020. NeuroCard: One Cardinality Estimator for All Tables. *Proc. VLDB Endow.* 14, 1 (2020), 61–73. https://doi.org/10.14778/3421424.3421432

[64] Zongheng Yang, Eric Liang, Amog Kamsetty, Chenggang Wu, Yan Duan, Xi Chen, Pieter Abbeel, Joseph M. Hellerstein, Sanjay Krishnan, and Ion Stoica. 2019. Deep Unsupervised Cardinality Estimation. *Proc. VLDB Endow.* 13, 3 (2019), 279–292. https://doi.org/10.14778/3368289.3368294

[65] Andreas Zeller. 2006. *Why programs fail - a guide to systematic debugging*. Elsevier.

[66] Zhuoyue Zhao, Robert Christensen, Feifei Li, Xiao Hu, and Ke Yi. 2018. Random Sampling over Joins Revisited. In *Proceedings of the 2018 International Conference on Management of Data, SIGMOD Conference 2018, Houston, TX, USA, June 10-15, 2018*, Gautam Das, Christopher M. Jermaine, and Philip A. Bernstein (Eds.). ACM, 1525–1539. https://doi.org/10.1145/3183713.3183739

[67] Rui Zhong, Yongheng Chen, Hong Hu, Hangfan Zhang, Wenke Lee, and Dinghao Wu. 2020. SQUIRREL: Testing Database Management Systems with Language Validity and Coverage Feedback. In *CCS '20: 2020 ACM SIGSAC Conference on Computer and Communications Security, Virtual Event, USA, November 9-13, 2020*, Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni Vigna (Eds.). ACM, 955–970. https://doi.org/10.1145/3372297.3417260

[68] Rong Zhu, Ziniu Wu, Yuxing Han, Kai Zeng, Andreas Pfadler, Zhengping Qian, Jingren Zhou, and Bin Cui. 2021. FLAT: Fast, Lightweight and Accurate Method for Cardinality Estimation. *Proc. VLDB Endow.* 14, 9 (2021), 1489–1502. https://doi.org/10.14778/3461535.3461539