Thanos: DBMS Bug Detection via Storage Engine Rotation Based Differential Testing

Ying Fu, Zhiyong Wu, Yuanliang Zhang, Jie Liang,

Jingzhou Fu, Yu Jiang *, Shanshan Li *, Xiangke Liao

国防科技大学
NATIONAL UNIVERSITY
OF DEFENSE TECHNOLOGY

清華大學
Tsinghua University

DBMS Differential Testing

- Establish test oracle efficiently

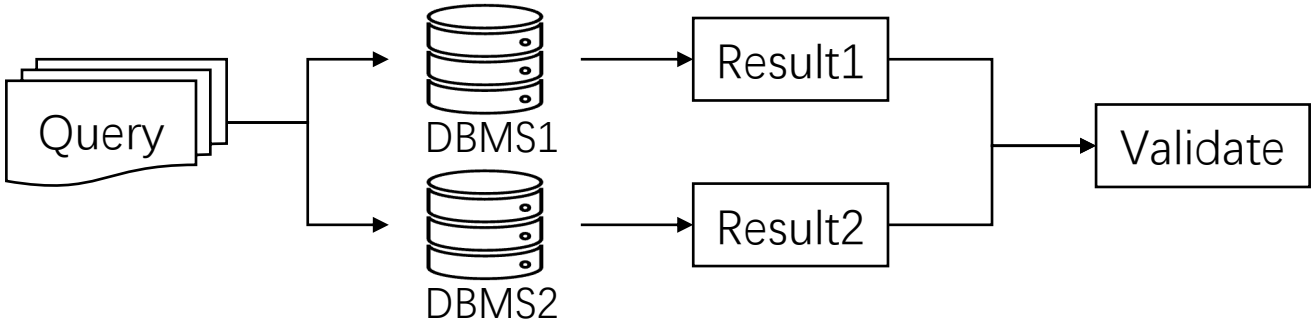

  - Version-Based Comparison: may overlook real issues due to minor version variability
  - Vendor-Based Comparison: limited by differences in system architecture and SQL compatibility issues

- Hard to balance tested DBMS diversity and input consistency

DBMS Storage Engine

- Diversity

  various storage engines with distinct features.

- Significance

  tied to core business functionalities

- Scalability

  New engines can be easily integrated

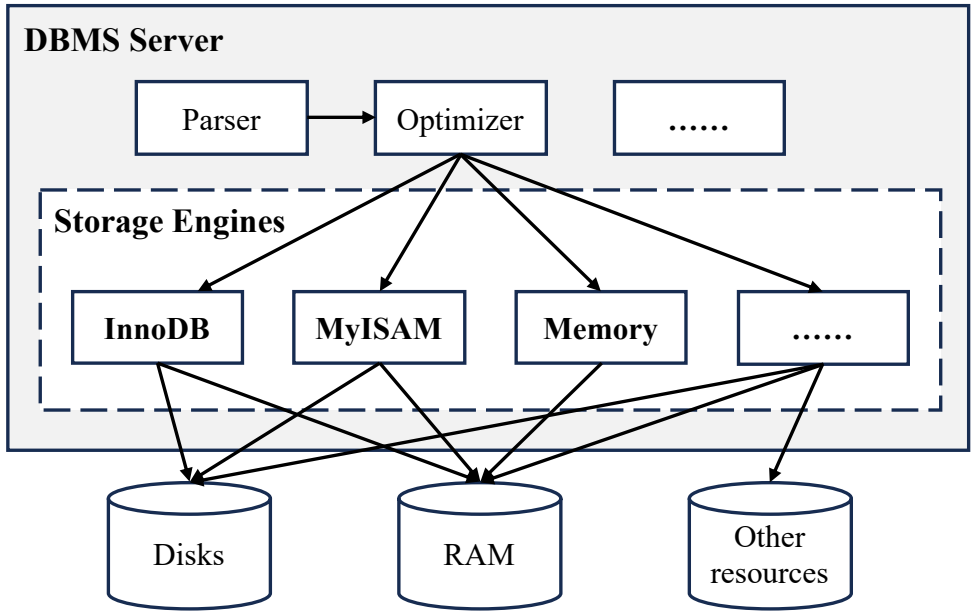

Switching storage engines -> Construct testable DBMS variants

Motivating Example

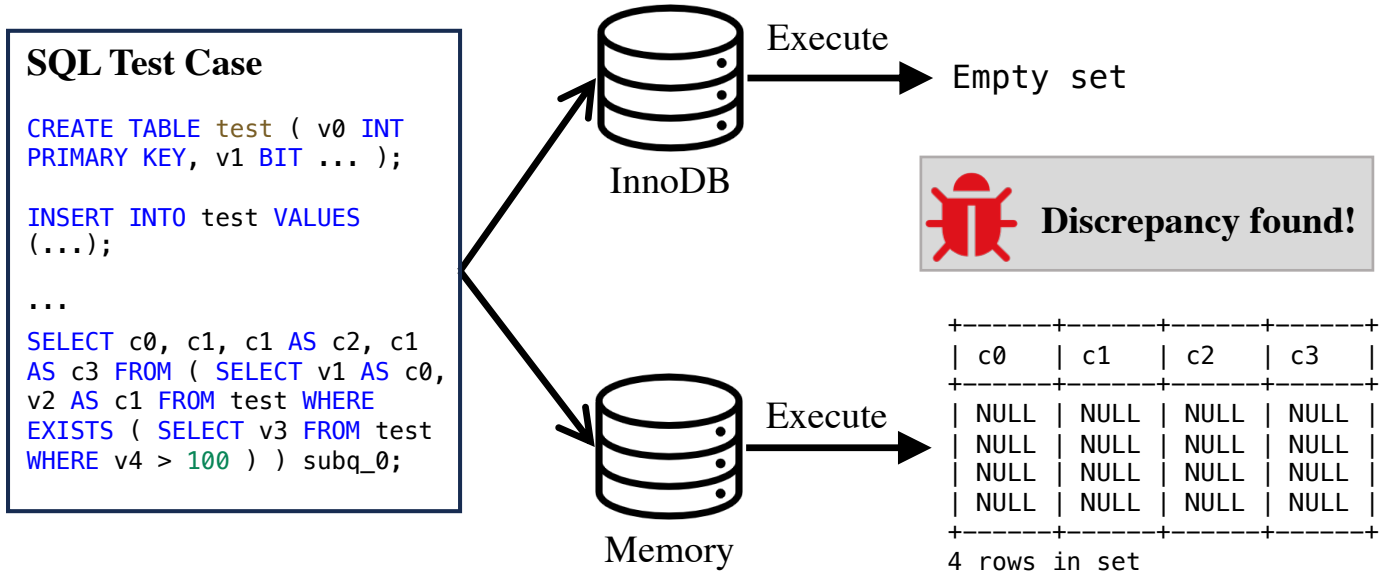

- Key Advantages: Blends implementation diversity with input-level consistency
- Key challenge: Ensure DBMS equivalence across diverse storage engines

Storage Engine Rotation Based Differential Testing

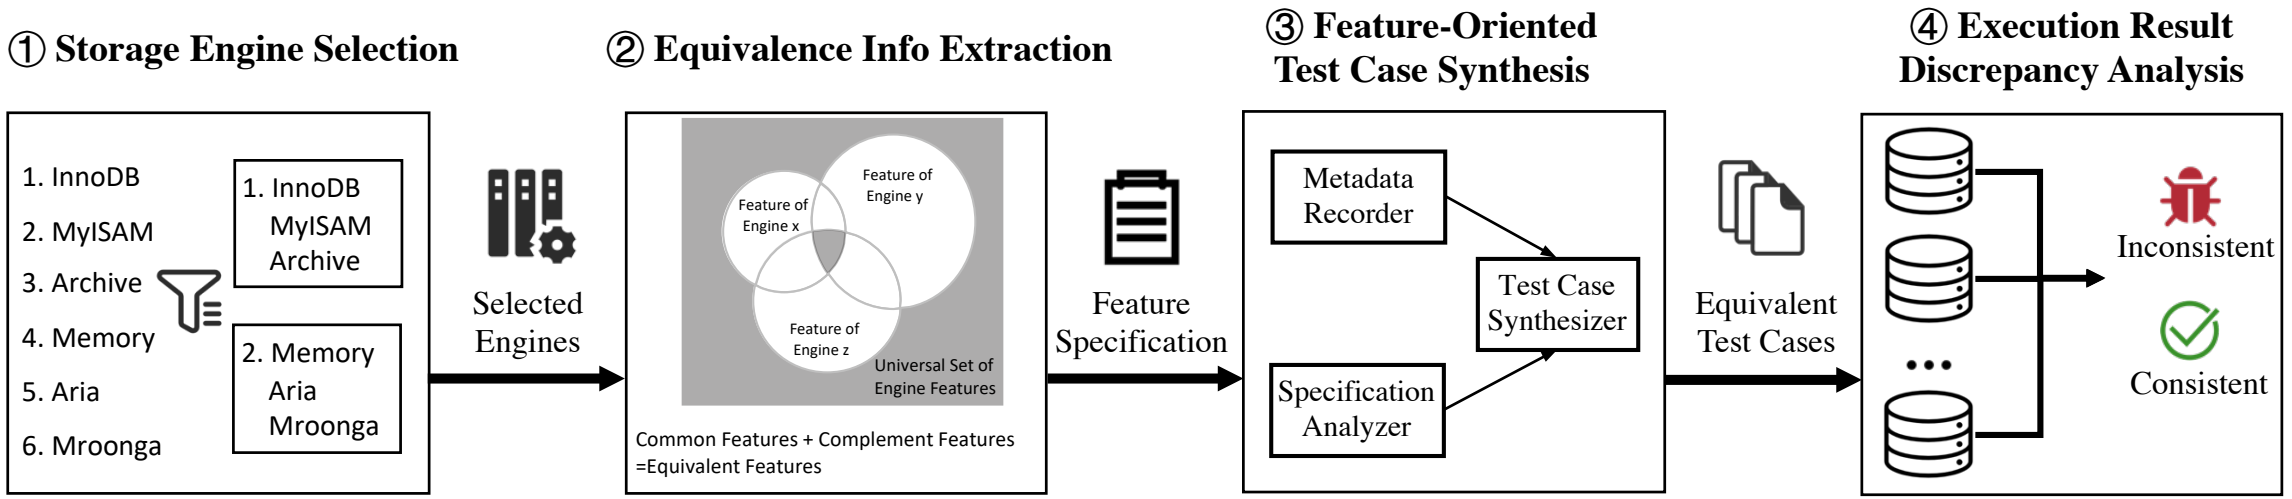

Feature of Storage Engine

- Specific functionalities and attributes
- Focus on SQL-visible features
- Classified into 10 major categories
  - Data type, Index Type, Data Integrity, Partition, Encryption, Compression, Transaction, Health-Check, Cache, Update statistics
  - For example: Index type includes B-Tree, Hash, Geospatial, Clustered, Multi-valued, Full-text
  - Each category encompasses approximately 2 to 40 specific features

	Data type			...	Index type		...	Health-check		Partition	Encryption	Transaction	...
	INT	BIT	BLOB		B-tree	Hash		Check	Repair				
InnoDB	✓	✓	✓		✓	×		✓	✓	✓	✓	✓	
MyISAM	✓	✓	✓		✓	×		✓	✓	✓	✓	×	
Archive	✓	×	✓		×	×		✓	✓	×	✓	×	
Mroonga	✓	✓	✓		✓	×		✓	✓	✓	✓	×	

Step ①&② Engine Selection & Equivalence Info Extraction

- Feature-Guided Storage Engine Selection
  - Focus on feature-rich combinations to uncover more bugs
  - Selecting impactful engine sets boosts testing effectiveness

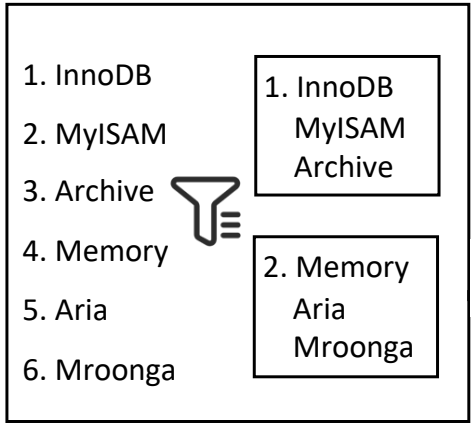

- Extract Equivalence Information

$$F_{equivalent} = \bigcap_{i=1}^{k} f_i \ \cup \ \left(U - \bigcup_{i=1}^{k} f_i \right)$$

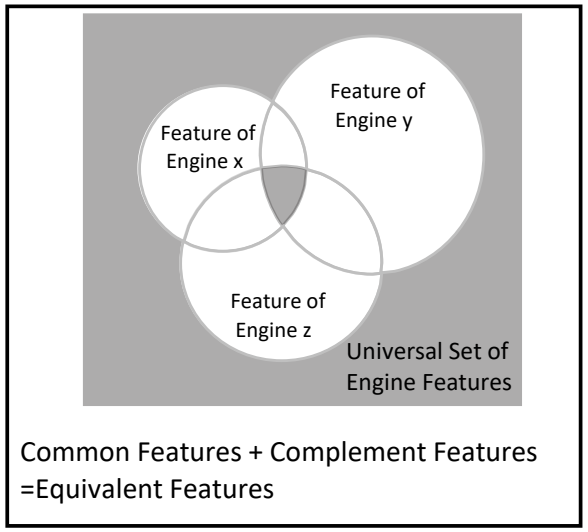

Step ③ Feature-Oriented Test Case Synthesis

- DDL Synthesis:

Create storage structures (e.g., CREATE TABLE, CREATE INDEX)

- DML/DQL Synthesis

Generate queries and data modification operations (e.g., INSERT, UPDATE, SELECT)

- Non-standard SQL Handling

Example: CHECK TABLE or ROW_FORMAT=COMPRESSED

- Metadata pool

Track table schemas, columns, and types

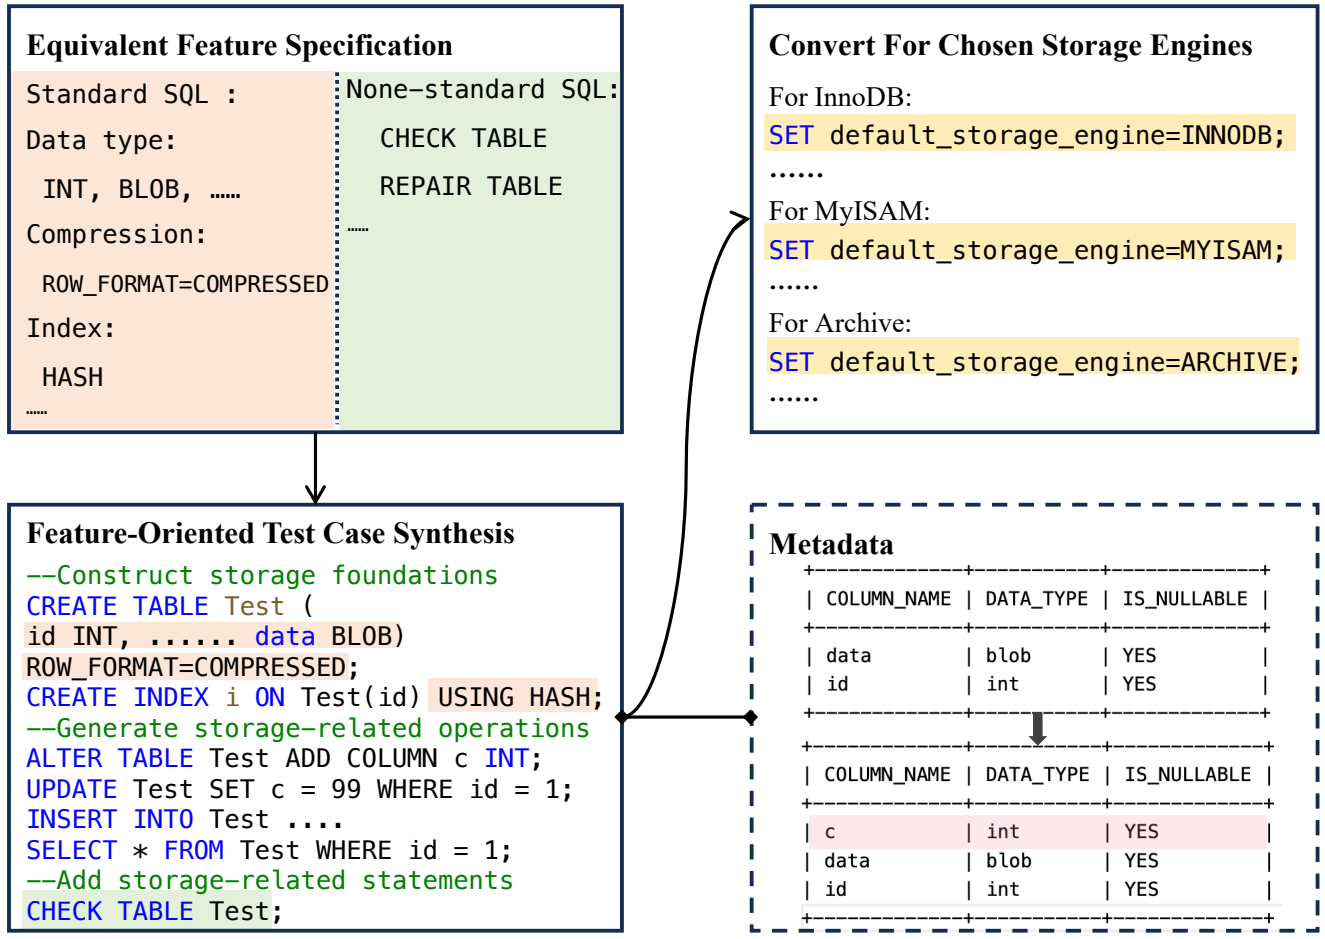

Step ④ Discrepancy Analysis of Execution Results

- Discrepancies in Query Results
- Discrepancies in Error Message
- Discrepancies in Database Final States
- Test case deduplication

Evaluation: New Bugs

- Evaluated on 3 DBMSs
  - MySQL、MariaDB、Percona
- 32 bugs
  - 29 were verified as Critical
  - Across 12 components

DBMS (Reported/Confirmed)	**Component**	**Bug Severity and Number**
MySQL (11/11)	Optimizer	Critical (2)
	DDL	Moderate (2)
	Options	Critical (1)
	Storage Engines	Critical (4), Moderate (2)
MariaDB (17/17)	Optimizer	Critical (7)
	Storage Engines	Critical (5)
	Window	Critical (1)
	Handler	Critical (1)
	Parser	Critical (2)
	Update	Critical (1)
Percona (4/4)	Optimizer	Critical (2)
	Storage Engines	Critical (1), Moderate (1)
Total	12 components	32 confirmed

Case Study: Hidden Bug in MariaDB

- Mroonga storage engine
- Discrepancies in database final states
- Incorrect handling the auto increment feature

```
SET default_storage_engine=Mroonga;
CREATE TABLE test(pk INT AUTO_INCREMENT, a INT, b
    INT, c INT, d INT, PRIMARY KEY (pk), KEY (a));
INSERT INTO test(a, b) VALUES (0,100),(200,2000);
INSERT INTO test(c, d) VALUES (0,100),(200,2000);
CREATE TRIGGER tr1 AFTER UPDATE ON test FOR EACH
    ROW SET @a= 100;
UPDATE test SET b = 3 WHERE a = 0;
-- Server crash
```

Evaluation: Comparison to Other Techniques

Bugs

DBMS	SQLancer	SQLsmith	SQUIRREL	THANOS
MySQL	0	1	1	6
MariaDB	0	0	1	5
Percona	0	0	1	3
Total	0	1	3	14
Increment	14 ↑	13 ↑	11 ↑	–

Covered Branches

DBMS	SQLancer	SQLsmith	SQUIRREL	THANOS
MySQL	59,242	93,742	109,323	120,156
MariaDB	60,293	88,923	100,920	132,532
Percona	63,829	89,987	109,823	143,293
Total	183,364	272,652	320,066	395,981
Increment	115.95% ↑	45.23% ↑	23.72% ↑	–

- More Bugs
- 24% ~ 116% more covered branches

Conclusion

- Idea: Constructs equivalent DBMS instances by switching storage engine
- Method: Feature-oriented test case synthesis
- Evaluation: 3 DBMSs, 32 bugs, 29/32 critical
- Future Work: Support single storage engine DBMSs

Thank you!

Please feel free to communicate via the following email.
fuying@nudt.edu.cn