

Data-Oriented Differential Testing of Object-Relational Mapping Systems

Thodoris Sotiropoulos, Stefanos Chaliasos, Vaggelis Atlidakis,
Dimitris Mitropoulos and Diomidis Spinellis

ICSE 2021
Distinguished Artifact Award

Object-Relational Mapping (ORM)

- Object-oriented interface on top of relational databases
- Promotes
 - Portability
 - Developers' productivity
- ORM frameworks are used by million of applications (e.g., OpenStack, Gitlab, Dropbox)

```
1 from django.db import models
2
3 class Person(models.Model):
4     age = models.IntegerField()
5     name = models.CharField(max_length=20)
6
7
8 p1 = Person(age=31, name="John")
9 p1.save()
10 p2 = Person.objects.get(age=32)
11 p2.delete()
```

```
INSERT INTO PERSON (age, name) VALUES (31, 'John')
```

```
SELECT * FROM PERSON WHERE AGE = 32 LIMIT 1
```

```
DELETE FROM PERSON WHERE ID = 2
```

ORM bugs (Django example)

```
1 q1 = T1.objects.using("mysql")
2 q2 = T2.objects.using("mysql")
3 q3 = T3.objects.using("mysql")
4
5 combined = q1.union(q2).union(q3)
6 for row in combined:
7     pass
```



```
(SELECT `t1`.`id` FROM `t1`)
UNION(
  (SELECT `t2`.`id` FROM `t2`)
  UNION(
    SELECT `t3`.`id` FROM `t3`))
```

```
django.db.utils.ProgrammingError: (1064, "You have an error in your SQL syntax;
check the manual that corresponds to your MySQL server version for the right syntax
to use near 'UNION (SELECT `listing`.`id`, `listing`.`foo` FROM `listing`)) subquery' at line 1")
```

Django generates a syntactically invalid SQL query with regards to MySQL

ORM bugs (peewee example)

```
1 expr = (1 + T.col)
2 squared = (expr * expr)
3 data = T.select(fn.sum(expr),
4               fn.avg(squared)).all()
5
6 for row in data:
7     print('avgExpr', row['avgExpr'])
```



```
SELECT SUM(1 + "t"."col"),
       AVG(1 + "t"."col" * 1 + "t"."col")
FROM "t" AS "t"
```

Expected SQL query

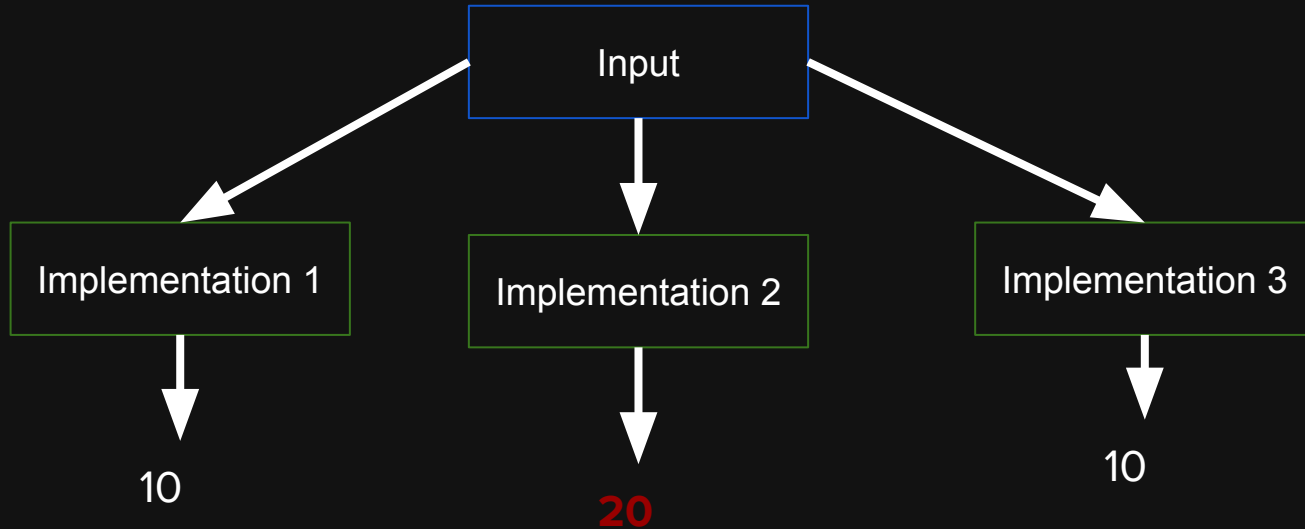
```
SELECT SUM(1 + "t"."col"),
       AVG((1 + "t"."col") * (1 + "t"."col"))
FROM "t" AS "t"
```

**Peewee generates a both syntactically and semantically valid SQL.
However, the query produces the wrong results.**

```
2 < avgExpr 19.00
3 < avgExpr 25.00
4 < avgExpr 29.00
5 < avgExpr 47.00
6 < avgExpr 7.00
7 < avgExpr 87.00
8 ---
9 > avgExpr 100.00
10 > avgExpr 16.00
11 > avgExpr 169.00
12 > avgExpr 1936.00
13 > avgExpr 225.00
14 > avgExpr 576.00
```

Test Oracle

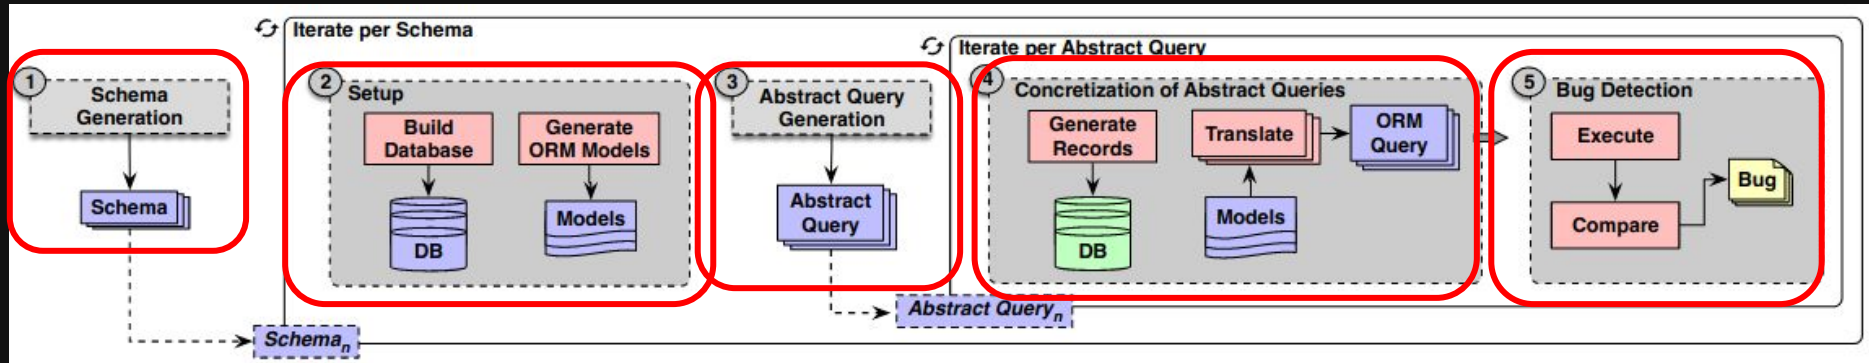
We use differential testing for establishing a test oracle



Challenges

- Lack of a common specification and input language
- Non-deterministic query results
- DBMS-dependent results (see Django bug)
- Data generation (see peewee bug)

Approach



1. Schema Generation
2. Schema Setup
3. Abstract Query Generation
4. Concretization of Abstract Queries
5. Bug Detection

Abstract Query Language (AQL)

- Supports wide range of operations (through functional notation)
 - Filtering
 - Sorting
 - Aliasing
 - Folding
 - Compound expressions
 - Aggregate functions
 - Unions / Intersections
- Closer to ORM APIs rather than SQL
- AQL queries are generated randomly up to a certain depth

```
q ∈ Query ::= eval qs | qs[i] | qs[i : i] | fold { (l :  $\alpha$  e)+ } qs  
qs ∈ QuerySet ::= new t | apply  $\lambda$  qs | qs ∪ qs  
| qs ∩ qs  
 $\lambda$  ∈ Func ::= filter p | map d | unique  $\phi$   
| sort ( $\phi$  asc) | sort ( $\phi$  desc)  
d ∈ FieldDecl ::= l : e | hidden l : e | d; d  
p ∈ Pred ::=  $\phi$  ⊕ e | p ∧ p | p ∨ p | ¬p  
e ∈ Expr ::= c |  $\phi$  |  $\alpha$  e | e + e | e − e | e * e | e / e  
 $\phi$  ∈ Field ::= t.c | l |  $\phi.c$   
 $\alpha$  ∈ AggrFunc ::= count | sum | avg | max | min  
⊕ ∈ BinaryOp ::= = | > | ≥ | < | ≤  
| contains | startswith | endswith
```


Data Generation

An AQL query is encoded as an SMT formula

```
1 apply(  
2   filter ("Table.str" contains "Paul")  
3   new Table  
4 )
```



A theorem prover generates assignments from which we derive executable INSERT statements

```
1 DELETE FROM "table";  
2 INSERT INTO "table"("id","str") VALUES (7,'Paul');  
3 INSERT INTO "table"("id","str") VALUES (-5,'!');  
4 INSERT INTO "table"("id","str") VALUES (-6,'H');  
5 INSERT INTO "table"("id","str") VALUES (13,'\xa0');  
6 INSERT INTO "table"("id","str") VALUES (0,'B');  
-
```



```
1 (declare-fun Table.id_1 () Int)  
2 (declare-fun Table.id_0 () Int)  
3 (declare-fun Table.id_2 () Int)  
4 (declare-fun Table.id_3 () Int)  
5 (declare-fun Table.id_4 () Int)  
6 (declare-fun Table.str_1 () String)  
7 (declare-fun Table.str_0 () String)  
8 (declare-fun Table.str_2 () String)  
9 (declare-fun Table.str_3 () String)  
10 (declare-fun Table.str_4 () String)  
11 (assert (and (not (= Table.id_3 Table.id_4))  
12   (not (= Table.id_2 Table.id_4))  
13   (not (= Table.id_2 Table.id_3))  
14   (not (= Table.id_1 Table.id_4))  
15   (not (= Table.id_1 Table.id_3))  
16   (not (= Table.id_1 Table.id_2))  
17   (not (= Table.id_0 Table.id_4))  
18   (not (= Table.id_0 Table.id_3))  
19   (not (= Table.id_0 Table.id_2))  
20   (not (= Table.id_0 Table.id_1))))  
21 (assert (and (not (= Table.str_3 Table.str_4))  
22   (not (= Table.str_2 Table.str_4))  
23   (not (= Table.str_2 Table.str_3))  
24   (not (= Table.str_1 Table.str_4))  
25   (not (= Table.str_1 Table.str_3))  
26   (not (= Table.str_1 Table.str_2))  
27   (not (= Table.str_0 Table.str_4))  
28   (not (= Table.str_0 Table.str_3))  
29   (not (= Table.str_0 Table.str_2))  
30   (not (= Table.str_0 Table.str_1))))  
31 (assert (or (str.suffixof "Paul" Table.str_0)  
32   (str.suffixof "Paul" Table.str_1)  
33   (str.suffixof "Paul" Table.str_2)  
34   (str.suffixof "Paul" Table.str_3)  
35   (str.suffixof "Paul" Table.str_4)))  
36
```

From AQL queries to ORM queries

- Use ORM-specific translators
- Each translator generates
 - The necessary boilerplate code (e.g., imports, db setup)
 - The actual ORM query
 - Code that prints results of the query

```
1 apply (  
2   filter "addCol" > 5  
3   apply (  
4     map "addCol": t1.colA + t1.t2.colB  
5     new t1  
6   )  
7 )  
-
```

```
1 import os, django  
2 from django.db.models import *  
3 os.environ.setdefault("DJANGO_SETTINGS_MODULE",  
4                       "djangoproject.settings")  
5 django.setup()  
6 from project.models import *  
7  
8 addCol = F("colA") + F("t2__colB")  
9 q = T1.objects.using("sqlite")\  
10   .annotate(addCol=addCol).filter(addCol__gt=5)\  
11   .values("addCol")  
12  
13 for r in q:  
14   print("addCol", r["addCol"])
```

Implementation Details

- We implement our approach as a tool called Cynthia
 - Implemented in Scala (~9k LoC)
 - Cynthia uses the Z3 theorem prover
- Cynthia currently provides support for five popular ORMs
 - Django
 - SQLAlchemy
 - Peewee
 - Sequelize
 - Activerecord
- ... and four DBMSs (Sqlite, MySQL, PostgreSQL, MS SQL Server)


Effectiveness

- Cynthia has found 28 bugs, of which 20 have been fixed.
- Most of the bugs have been discovered in Django and SQLAlchemy
- DBMS-dependent bugs (**11 / 28**)
- Most of DBMS-dependent bugs are triggered when the code is run on top of PostgreSQL and MSSQL

ORM	Total	Fixed	Confirmed	Unconfirmed
Django	10	6	3	1
SQLAlchemy	8	8	0	0
Sequelize	5	2	1	2
peewee	4	4	0	0
Activerecord	1	0	1	0
Total	28	20	5	3

Effectiveness of Solver-Based Data Generation

- We compared our solver-based approach with a “naive” approach
 - I.e., generating random records without considering the constraints of the generated queries
- We spawned 20 testing sessions consisting of 100 AQL queries, and measured in how many queries the ORMs returned empty results
- Unsatisfied Queries (Solver-based approach): **7.9%**
- Unsatisfied Queries (“Naive” approach): **38%**
- We get no improvement even if we generate more records
 - generating 50 random records is the same with generating 1000 random records



Conclusion

- Introduced the first data-oriented differential testing approach for systematically testing ORM implementations
- We showed that differential testing can be also applicable to (seemingly) dissimilar interfaces, such as ORMs
- We showed that compared with other simplistic approaches, our solver-based approach enhances the bug detection capability, and is suitable for differential testing
- Our tool, Cynthia, discovered 28 bugs, most of which have been fixed by the developers.
- The effectiveness of Cynthia can be improved by considering
 - other forms of queries (e.g., write queries)
 - transaction management

Thank you



Tool: <https://github.com/theosotr/cynthia>

Artifact: <https://doi.org/10.5281/zenodo.4455486>

Characteristics of Discovered Bugs

Type	# Bugs	All	SQLite	MySQL	PostgreSQL	MSSQL
Logic Error	12	11	0	0	0	1
Invalid SQL	11	3	1	3	2	3
Crash	5	3	0	0	2	0
Total	28	17	1	3	4	4

- Most of the discovered bugs are logic errors (**12 / 28**)
- Followed by “Invalid SQL” bugs (**11 / 28**) and crashes (**5 / 28**)
- Almost all “logic errors” are DBMS-independent
- Yet, there is a large number of DBMS-dependent bugs (**11 / 28**)
- Most of DBMS-dependent bugs are triggered when the code is run on top of PostgreSQL and MSSQL

Bug Detection

- We make DBMS-specific comparisons
- A bug is found when one of the following holds
 - Two ORMs produce different results on the same DBMS.
 - An ORM query is successfully run on a specific DBMS, but the same query written in another ORM fails on the same DBMS.

Concretization of Abstract Queries

- Data Generation
- Translation of AQL query into a concrete ORM query