

Automated Reasoning of Database Queries

Shumo Chu, University of Washington

with Konstantin Weitz, Chenglong Wang, Daniel Li, Brendan Murphy,
Jared Roesch, Alvin Cheung, Dan Suciu

- SQL: a language supported by all relational databases
- Restricted abstraction enabling powerful optimizations
- 30 years of research results in many optimizations based on semantic equivalent rewrites



<https://insights.stackoverflow.com/survey/2017>

Lack tools that can reason about SQL equivalences

Automated Solver for SQL:

$Q1 = Q2?$



Automated Solver for SQL:

∀ possible input, $Q1 = Q2$?



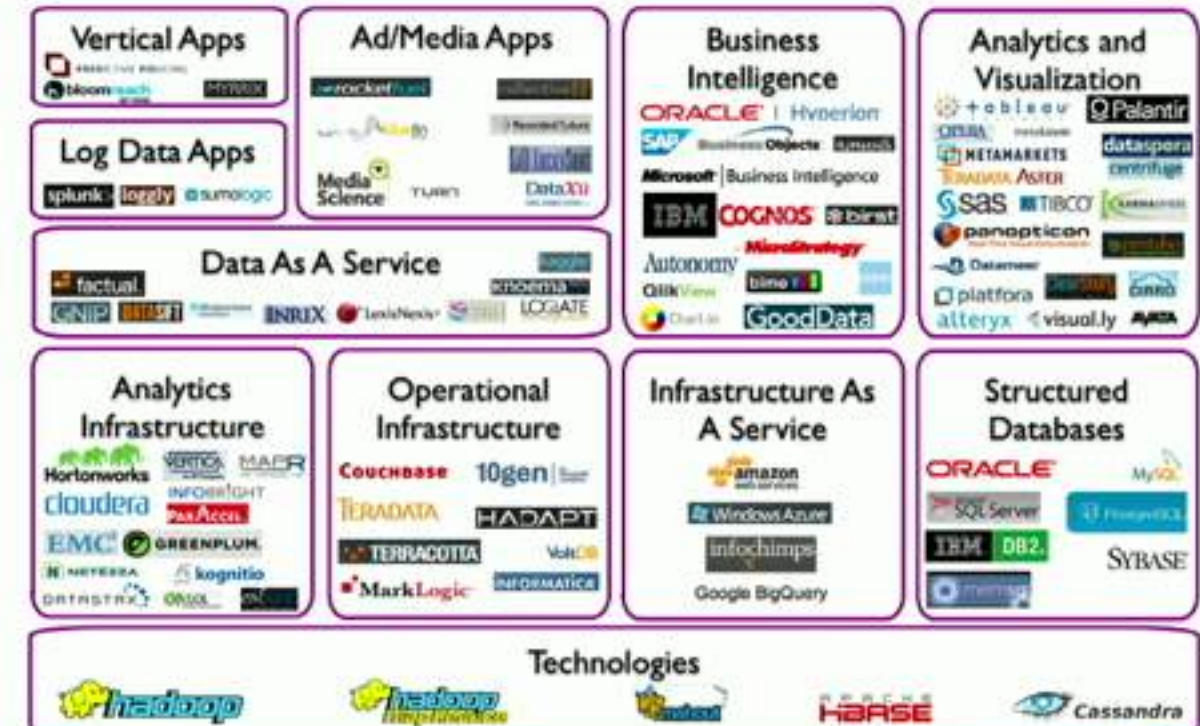
Automated Solver for SQL:

▽ possible input, $Q1 = Q2?$

- Correctness of query rewrite



Big Data Landscape



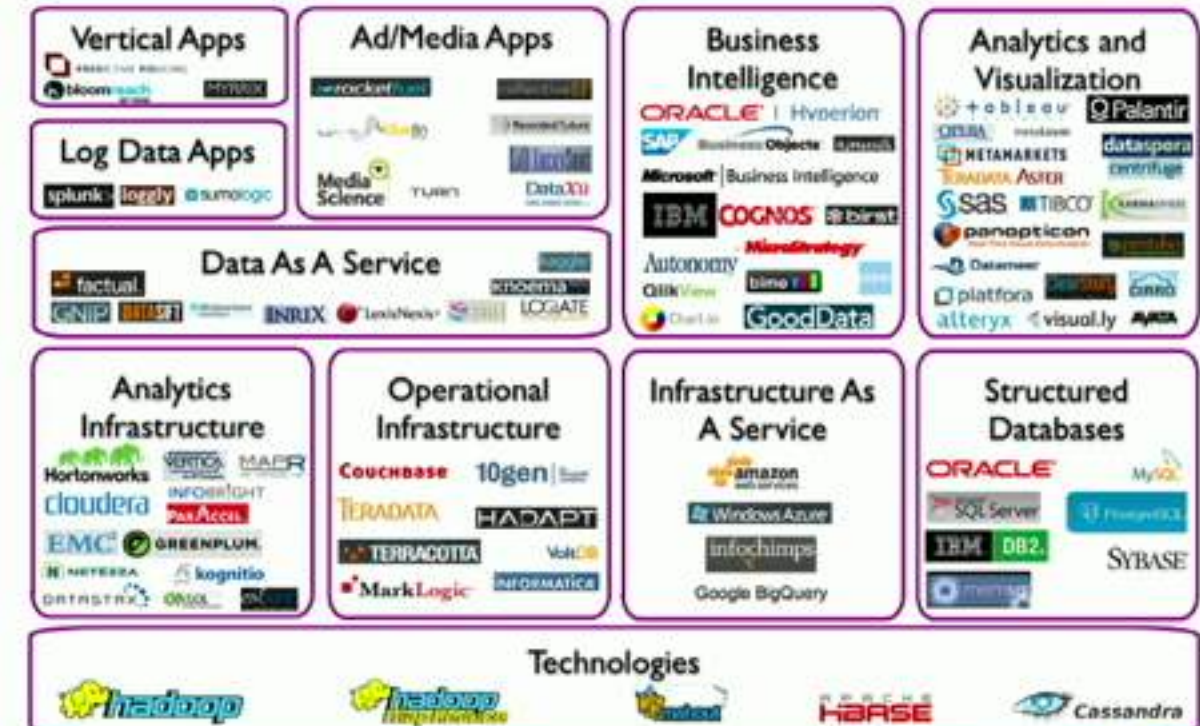
Automated Solver for SQL:

▽ possible input, $Q1 = Q2?$

- Correctness of query rewrite
- Semantics layer of data systems



Big Data Landscape



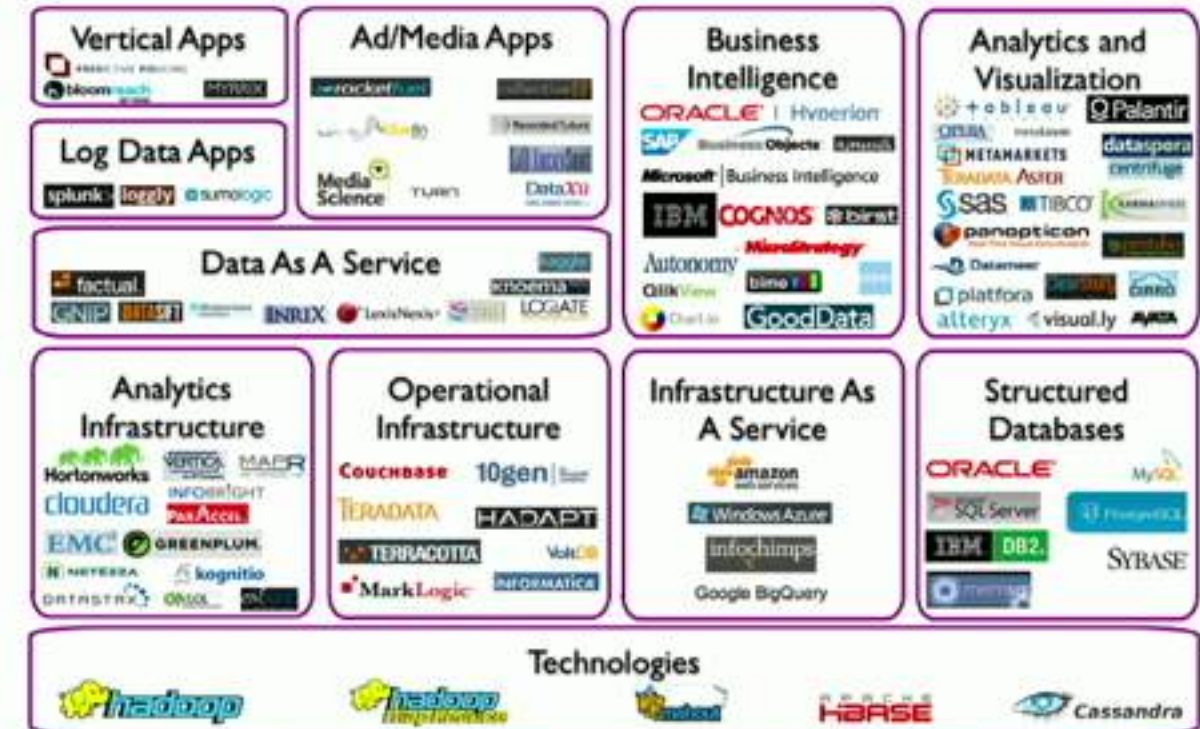
Automated Solver for SQL:

▽ possible input, $Q1 = Q2?$

- Correctness of query rewrite
- Semantics layer of data systems
- Auto grading on SQL assignments



Big Data Landscape



Challenges

- Checking equivalences of two FO sentences over finite models is undecidable ☹
- Rich language features
 - Aggregation and grouping
 - Index and integrity constraints
 - Correlated subqueries



Boris A.
Trakhtenbrot

Undecidability \neq No Proof



Interactive Theorem Prover that
validates mechanized proofs

Undecidability \neq No Proof



Interactive Theorem Prover that
validates mechanized proofs

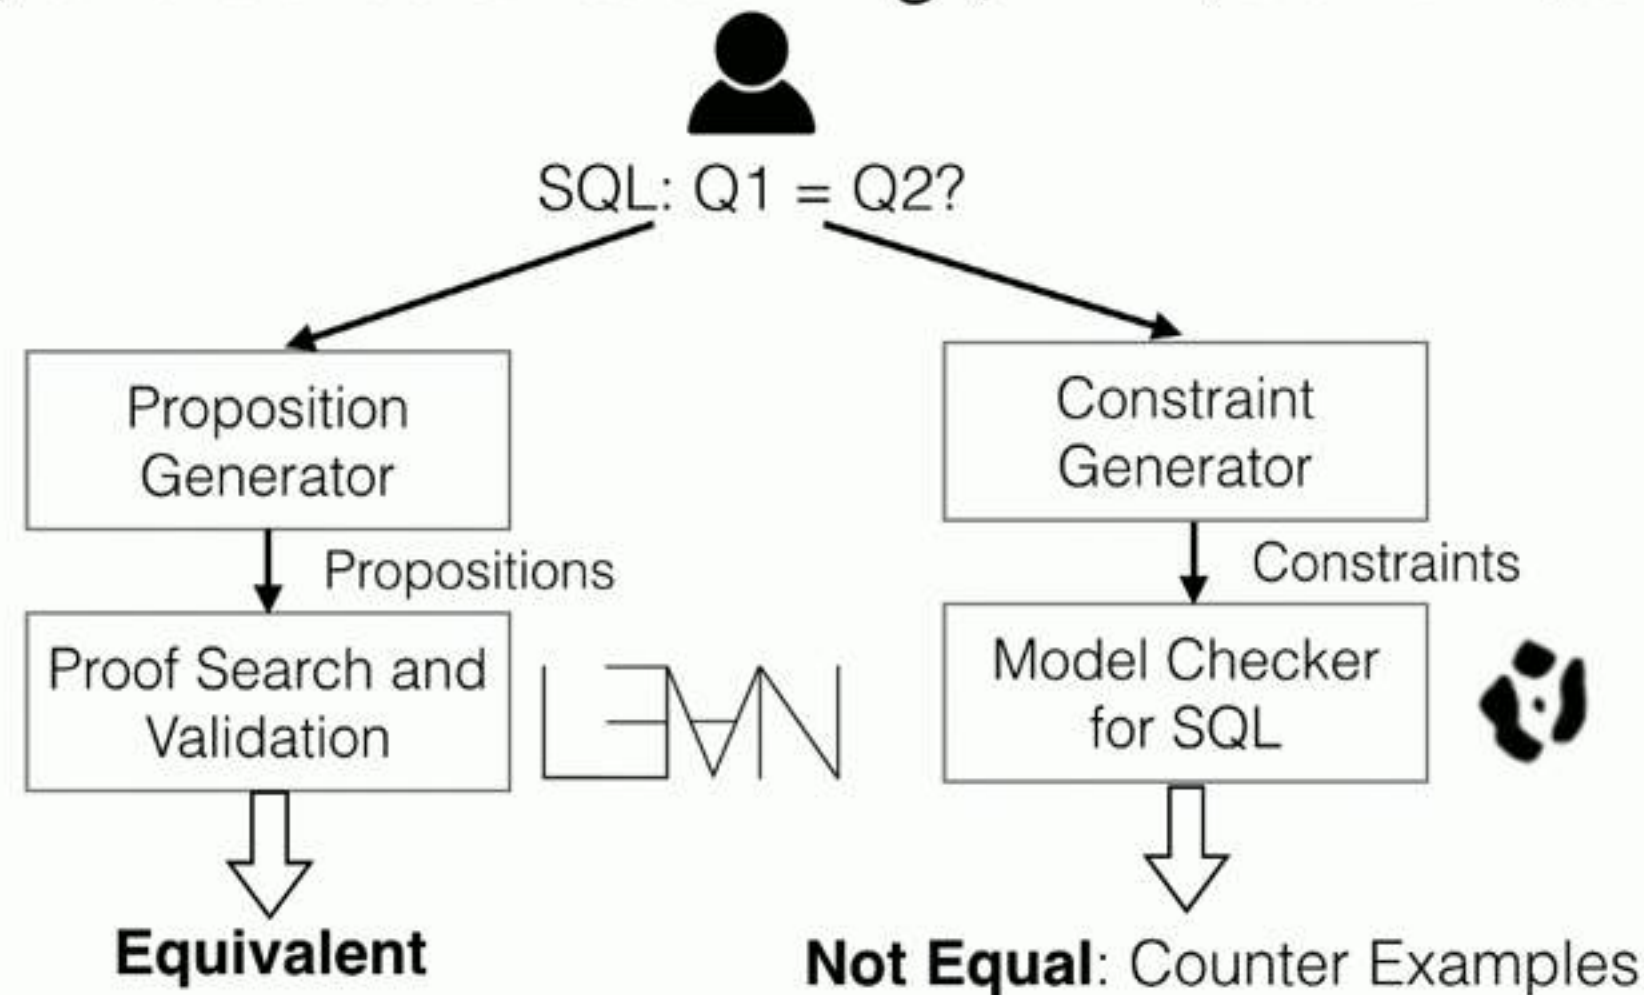
**The model of an inequivalence
is usually not large**



Constraint Solver for Model
Checking

Construct a SQL Solver

- Cosette: An (almost) automated solver for SQL by *combining interactive theorem proving and constraints solving* (PLDI 17, SIGMOD17, CIDR17, VLDB18)



The Search for Formal SQL Semantics

What is SQL?

ISO/IEC 9075-2:2016 [Preview](#)

Information technology -- Database languages -- SQL -- Part 2: Foundation (SQL/Foundation)

ISO/IEC 9075-2:2016 defines the data structures and basic operations on SQL-data. It provides functional capabilities for creating, accessing, maintaining, controlling, and protecting SQL-data.

ISO/IEC 9075-2:2016 specifies the syntax and semantics of a database language:

- For specifying and modifying the structure and the integrity constraints of SQL-data.
- For declaring and invoking operations on SQL-data and cursors.
- For declaring database language procedures.
- For embedding SQL-statements in a compilation unit that is otherwise written in a particular programming language (host language).
- For deriving an equivalent compilation unit in the host language. In that equivalent compilation unit, each which invoke an SQL externally-invoked procedure that, when executed, has an effect equivalent to executing the SQL-statement.
- For direct invocation of SQL-statements.
- To support dynamic preparation and execution of SQL-statements.

ISO/IEC 9075-2:2016 provides a vehicle for portability of data definitions and compilation units between SQL-implementations.

ISO/IEC 9075-2:2016 provides a vehicle for interconnection of SQL-implementations.

Implementations of this ISO/IEC 9075-2:2016 can exist in environments that also support application programming languages, end-user query languages, report generator systems, data dictionary systems, program library systems, and distributed communication systems, as well as various tools for database design, data administration, and performance optimization.

Buy this standard

Format

Language



PDF

English

CHF 198

Buy

General information

Status : Published

Publication date : 2016-12

Edition : 5

Number of pages : 1707

Technical Committee : ISO/IEC JTC 1/SC 32 Data management and interchange

ICS : 35.060 Languages used in information technology

- More than **1,700** pages natural language description
- Insufficient for a rigorous SQL semantics

What is SQL?

- Data model: relations
 - A relation is an unordered collection of tuples
- Schema: set of attributes, each tuple in a relation must have same schema
- SQL query: $Q(R_1, R_2, \dots, R_k) = R$

Relational Algebra

Relational Algebra (\bowtie , σ , Π , \cup , $-$)

- \bowtie : e.g. $Courses \bowtie Student = \text{SELECT } * \text{ FROM } Course, Student$
- σ : e.g. $\sigma_{len > 100}(Movie) = \text{SELECT } * \text{ FROM } Movie \text{ WHERE } len > 100$
- Π : e.g. $\Pi_{first}(Name) = \text{SELECT } first \text{ FROM } Name$
- \cup : e.g. $R \cup S = (\text{SELECT } * \text{ FROM } R) \text{ UNION ALL } (\text{SELECT } * \text{ FROM } S)$

Relational Algebra

Relational Algebra ($\bowtie, \sigma, \Pi, \cup, -$) UCQ

- \bowtie : e.g. $Courses \bowtie Student = \text{SELECT } * \text{ FROM } Course, Student$
- σ : e.g. $\sigma_{len > 100}(Movie) = \text{SELECT } * \text{ FROM } Movie \text{ WHERE } len > 100$
- Π : e.g. $\Pi_{first}(Name) = \text{SELECT } first \text{ FROM } Name$
- \cup : e.g. $R \cup S = (\text{SELECT } * \text{ FROM } R) \text{ UNION ALL } (\text{SELECT } * \text{ FROM } S)$

Relational Algebra

- More operators needed
 - δ : e.g. $\delta(R) = \text{SELECT DISTINCT } * \text{ FROM } R$
 - γ : e.g. $\gamma_{\text{name, count}(*)} \text{Person} = \text{SELECT name, count}(*) \text{ FROM Person GROUP BY name}$
- Constraints: keys, foreign keys ...
- Lacking “meta-theory”

Relational Algebra

- More operators needed
 - δ : e.g. $\delta(R) = \text{SELECT DISTINCT } * \text{ FROM } R$
 - γ : e.g. $\gamma_{\text{name, count}(*)} \text{Person} = \text{SELECT name, count}(*) \text{ FROM Person GROUP BY name}$
- Constraints: keys, foreign keys ...
- Lacking “meta-theory”



Leaking Abstraction: What is the semantics of RA?

Relation as List

Existing Coq Formalization [Malecha et al., POPL 10]

- Two queries are equivalent if returning the **same** result for **all possible** input relations
- Need to reason about two lists are equivalent up to permutations

Q1 = **SELECT ***
FROM (R UNION ALL S)
WHERE b

Q2 = (**SELECT * FROM R WHERE b**)
UNION ALL
(**SELECT * FROM S WHERE b**)

Q1 = **SELECT ***
FROM (R UNION ALL S)
WHERE b

Q2 = (**SELECT * FROM R WHERE b**)
UNION ALL
(**SELECT * FROM S WHERE b**)

Q1 = Q2 ?

Induction on R:


Assume Q1 == Q2 when R has N tuples
Then when R is of size N+1:

...

Induction on S:

Assume Q1 == Q2 when S has N tuples
Then when S is of size N+1:

...



Q1 = **SELECT ***
FROM (R UNION ALL S)
WHERE b

Q2 = (**SELECT * FROM R WHERE b**)
UNION ALL
(**SELECT * FROM S WHERE b**)

Q1 = Q2 ?

Induction on R:


Assume Q1 == Q2 when R has N tuples
Then when R is of size N+1:

...

Induction on S:

Assume Q1 == Q2 when S has N tuples
Then when S is of size N+1:

...



**400 Line of Coq for a simple rewrite, very limited
rewrites has been proven**

Univalent SQL Semantics

[Chu et al., PLDI 17]

Relation as Function

$R: \text{Relation} \rightarrow \llbracket R \rrbracket: \text{Tuple} \rightarrow \mathbb{N}$

Relation as Function

$R:\text{Relation} \rightarrow \llbracket R \rrbracket:\text{Tuple} \rightarrow \mathbb{N}$

$b:\text{Predicate} \rightarrow \llbracket b \rrbracket:\text{Tuple} \rightarrow \{0, 1\}$

$\text{SELECT } * \text{ FROM } R \text{ WHERE } b \rightarrow Q(t) = \llbracket R \rrbracket t \times \llbracket b \rrbracket t$

$R \text{ UNION ALL } S \rightarrow Q(t) = \llbracket R \rrbracket t + \llbracket S \rrbracket t$

Q1 = **SELECT ***
FROM (R UNION ALL S)
WHERE b

Q2 = (**SELECT * FROM R WHERE b**)
UNION ALL
(**SELECT * FROM S WHERE b**)

Q1 = **SELECT ***
FROM (R UNION ALL S)
WHERE b



$Q1(t) = ([R](t) + [S](t)) \times [b](t)$

Q2 = (**SELECT * FROM R WHERE b**)
UNION ALL
(**SELECT * FROM S WHERE b**)



$Q2(t): [R](t) \times [b](t) + [S](t) \times [b](t)$

Q1 = **SELECT** *
FROM (R **UNION ALL** S)
WHERE b



$Q1(t) = ([R](t) + [S](t)) \times [b](t)$


Q2 = (**SELECT** * **FROM** R **WHERE** b)
UNION ALL
(**SELECT** * **FROM** S **WHERE** b)



$Q2(t): [R](t) \times [b](t) + [S](t) \times [b](t)$

Q1 = Q2 ?

Algebraic Reasoning



QED

Problem 1: Projection

[[**SELECT** first **FROM** Name]]= ?



Problem 1: Projection

[[**SELECT** first **FROM** Name]]= ?

Name

First	Last
Michael	Shulman
Michael	Jordan
Alex	Aiken
Alex	Krizhevsky



First
Michael
Michael
Alex
Alex

Add up multiplicities:
 $Q(\text{"Michael"}) = 2$
 $Q(\text{"Alex"}) = 2$

Problem 1: Projection

`[[SELECT first FROM Name]] = ?`

$$Q(t') = \sum_{t: \text{Tuple}} [t.\text{first} = t'.\text{first}] \times [[\text{Name}]](t)$$

Problem 1: Projection

$\llbracket \text{SELECT first FROM Name} \rrbracket = ?$

0 or 1

$$Q(t') = \sum_{t: \text{Tuple}} [t.\text{first} = t'.\text{first}] \times \llbracket \text{Name} \rrbracket(t)$$

Issues:

- Sum can potentially be infinite
- Need to reason about equality with algebraic expressions (now with sums ...)

Problem 2: Duplicate Elimination

[[SELECT DISTINCT first FROM Name]] = ?

Proposal: A step functions over \mathbb{N} : e.g. $S(0) = 0$, $S(x) = 1$

Issues: Back to inductive proof again 🙄

Univalent Semantics

- HoTT Type instead of \mathbb{N}
- Dependent Pair Type (Σ -Type) to represent summation
- Squash Type to represent duplicate elimination
- **Bonus: Unifies types and propositions**



P. Seshadri, J. Hellerstein, H. Pirahesh, T. Y. Leung,
R. Ramakrishnan, D. Srivastava, P. Stuckey, S. Sudarshan

Cost-Based Optimization for Magic: Algebra and Implementation.
SIGMOD 1996

Introduction of θ -semijoin:

$$R_1 \bowtie_{\theta} R_2 \equiv R_1 \bowtie_{\theta} (R_2 \ltimes_{\theta} R_1)$$

Pushing θ -semijoin through join:

$$(R_1 \bowtie_{\theta_1} R_2) \ltimes_{\theta_2} R_3 \equiv (R_1 \bowtie_{\theta_1} R'_2) \ltimes_{\theta_2} R_3$$

$$R'_2 = E_2 \ltimes_{\theta_1 \wedge \theta_2} (R_1 \bowtie R_3)$$

Pushing θ -semijoin through aggregation:

$$\bar{g} \mathcal{F}_{\bar{f}}(R_1) \ltimes_{c_1=c_2} R_2 \equiv_{\bar{g}} \mathcal{F}_{\bar{f}}(R_1 \ltimes_{c_1=c_2} R_2)$$



P. Seshadri, J. Hellerstein, H. Pirahesh, T. Y. Leung,
R. Ramakrishnan, D. Srivastava, P. Stuckey, S. Sudarshan

Cost-Based Optimization for Magic: Algebra and Implementation.
SIGMOD 1996

Introduction of θ -semijoin:

$$R_1 \bowtie_{\theta} R_2 \equiv R_1 \bowtie_{\theta} (R_2 \ltimes_{\theta} R_1)$$



Pushing θ -semijoin through join:

$$(R_1 \bowtie_{\theta_1} R_2) \ltimes_{\theta_2} R_3 \equiv (R_1 \bowtie_{\theta_1} R'_2) \ltimes_{\theta_2} R_3$$

$$R'_2 = E_2 \ltimes_{\theta_1 \wedge \theta_2} (R_1 \bowtie R_3)$$



Pushing θ -semijoin through aggregation:

$$\bar{g} \mathcal{F}_{\bar{f}}(R_1) \ltimes_{c_1=c_2} R_2 \equiv_{\bar{g}} \mathcal{F}_{\bar{f}}(R_1 \ltimes_{c_1=c_2} R_2)$$



Can we do better?

Axiomatic Foundations and Semi-Decision Procedures

[Chu et al., VLDB 18]

Idea 1: projection requires summation

$$\begin{aligned} \llbracket \textbf{SELECT first FROM Name} \rrbracket &\equiv \\ \llbracket Q \rrbracket(t') &= \sum_t [t.\text{name} = t'.\text{name}] \times \llbracket R \rrbracket(t) \end{aligned}$$

Solution: new* infinitary operator \sum_t

Idea 2: **DISTINCT** converts bags to sets

$$\begin{aligned} & \llbracket \text{SELECT DISTINCT first FROM Name} \rrbracket \equiv \\ \llbracket Q \rrbracket(t') &= \quad \llbracket \sum_t [t.\text{first} = t'.\text{first}] \times \llbracket \text{Name} \rrbracket(t) \rrbracket \end{aligned}$$

Solution: add the squash^{*} operator $\llbracket \dots \rrbracket$

Intuition: $\llbracket 0 \rrbracket = 0$ and $\llbracket x \rrbracket = 1$ otherwise

Idea 3: non-monotone operators e.g. **EXCEPT, NOT EXIST**

Solution: add a **not(...)** operator

Intuition: $\text{not}(\emptyset) = 1$, $\text{not}(x) = \emptyset$ otherwise

U-Semirings

Definition: An **Unbounded-semiring** is a structure $(\mathbf{U}, 0, 1, +, \times, \|\dots\|, \text{not}, (\sum_{\mathbf{D}})_{\mathbf{D} \in \mathcal{D}})$ satisfying a list of axioms

U-Semirings

Comm. Semiring

Definition: An **Unbounded-semiring** is a structure $(\mathbf{U}, 0, 1, +, \times, \|\dots\|, \text{not}, (\sum_{\mathbf{D}})_{\mathbf{D} \in \mathcal{D}})$ satisfying a list of axioms

Example: Axiomatization of Sum

$$\sum_t (f_1(t) + f_2(t)) = \sum_t f_1(t) + \sum_t f_2(t)$$

Distributivity

$$\sum_{t_1} \sum_{t_2} f(t_1, t_2) = \sum_{t_2} \sum_{t_1} f(t_1, t_2)$$

Commutativity

$$x \times \sum_t f(t) = \sum_t x \times f(t)$$

Associativity

$$\| \sum_t f(t) \| = \| \sum_t \|f(t)\| \|$$

Idempoten on Squash

Axiomization of Integrity Constraints

- The **key constraint** on R.k is the identity:

$$[t.k=t'.k] \times R(t) \times R(t') = [t=t'] \times R(t)$$

Axiomization of Integrity Constraints

- The **key constraint** on $R.k$ is the identity:

$$[t.k=t'.k] \times R(t) \times R(t') = [t=t'] \times R(t)$$

Axiomization of Integrity Constraints

- The **key constraint** on R.k is the identity:

$$[t.k=t'.k] \times R(t) \times R(t') = [t=t'] \times R(t)$$

- The **foreign key constraint** from S.fk to R.k is :

$$S(t') = S(t') \times \sum_t R(t) \times [t.k=t'.fk]$$

Proving Equivalences with ICs

```
-- Q1
SELECT ip.np, itm.type, itm.itemno
FROM (SELECT DISTINCT itp.itemno AS itn,
                    itp.np AS np
      FROM price WHERE price.np > 1000) ip, itm
WHERE ip.itn = itm.itemno;

-- Q2
SELECT DISTINCT price.np, itm.type, itm.itemno
FROM price, itm
WHERE price.np > 1000 AND itp.itemno = itm.itemno;
```

itemno is a key of itm

Proving Equivalences with ICs

$$Q_1(t) = \sum_{t_1, t_2} [t_1.np = t.np] \times [t_2.type = t.type] \times [t_2.itemno = t.itemno] \times [t_1.itn = t_2.itemno] \times \parallel \sum_{t'} [t'.itemno = t_1.itn] \times [t'.np = t_1.np] \times [t'.np > 1000] \times price(t') \parallel \times itm(t_2)$$

$$Q_2(t) = \parallel \sum_{t_1, t_2} [t_2.type = t.type] \times [t_2.itemno = t.itemno] \times [t_1.itemno = t.itemno] \times [t_2.itemno = t_1.itemno] \times [t_1.np = t.np] \times [t_1.np > 1000] \times price(t_1) \times itm(t_2) \parallel$$

Proving Equivalences with ICs

$$Q_1(t) = \sum_{t_1, t_2} [t_1.np = t.np] \times [t_2.type = t.type] \times [t_2.itemno = t.itemno] \times [t_1.itn = t_2.itemno] \times \parallel \sum_{t'} [t'.itemno = t_1.itn] \times [t'.np = t_1.np] \times [t'.np > 1000] \times price(t') \parallel \times itm(t_2)$$

$$Q_1(t) = \parallel \sum_{t_2, t'} [t_2.type = t.type] \times [t_2.itemno = t.itemno] \times [t'.itemno = t.itemno] \times [t'.np = t.np] \times [t'.np > 1000] \times price(t') \times itm(t_2) \parallel$$

$$Q_2(t) = \parallel \sum_{t_1, t_2} [t_2.type = t.type] \times [t_2.itemno = t.itemno] \times [t_1.itemno = t.itemno] \times [t_2.itemno = t_1.itemno] \times [t_1.np = t.np] \times [t_1.np > 1000] \times price(t_1) \times itm(t_2) \parallel$$



**Rewrite Using
U-Semiring Axioms**

Algorithm: Main Idea

Check $\llbracket Q_1 \rrbracket = \llbracket Q_2 \rrbracket$ by generalizing two known cases:

- UCQ under set semantics:
 - Check for homomorphisms $Q_1 \leftrightarrow Q_2$
- UCQ under bag semantics:
 - Check for isomorphisms $Q_1 \rightarrow Q_2$
- “Chase” the axiomatization of constraints

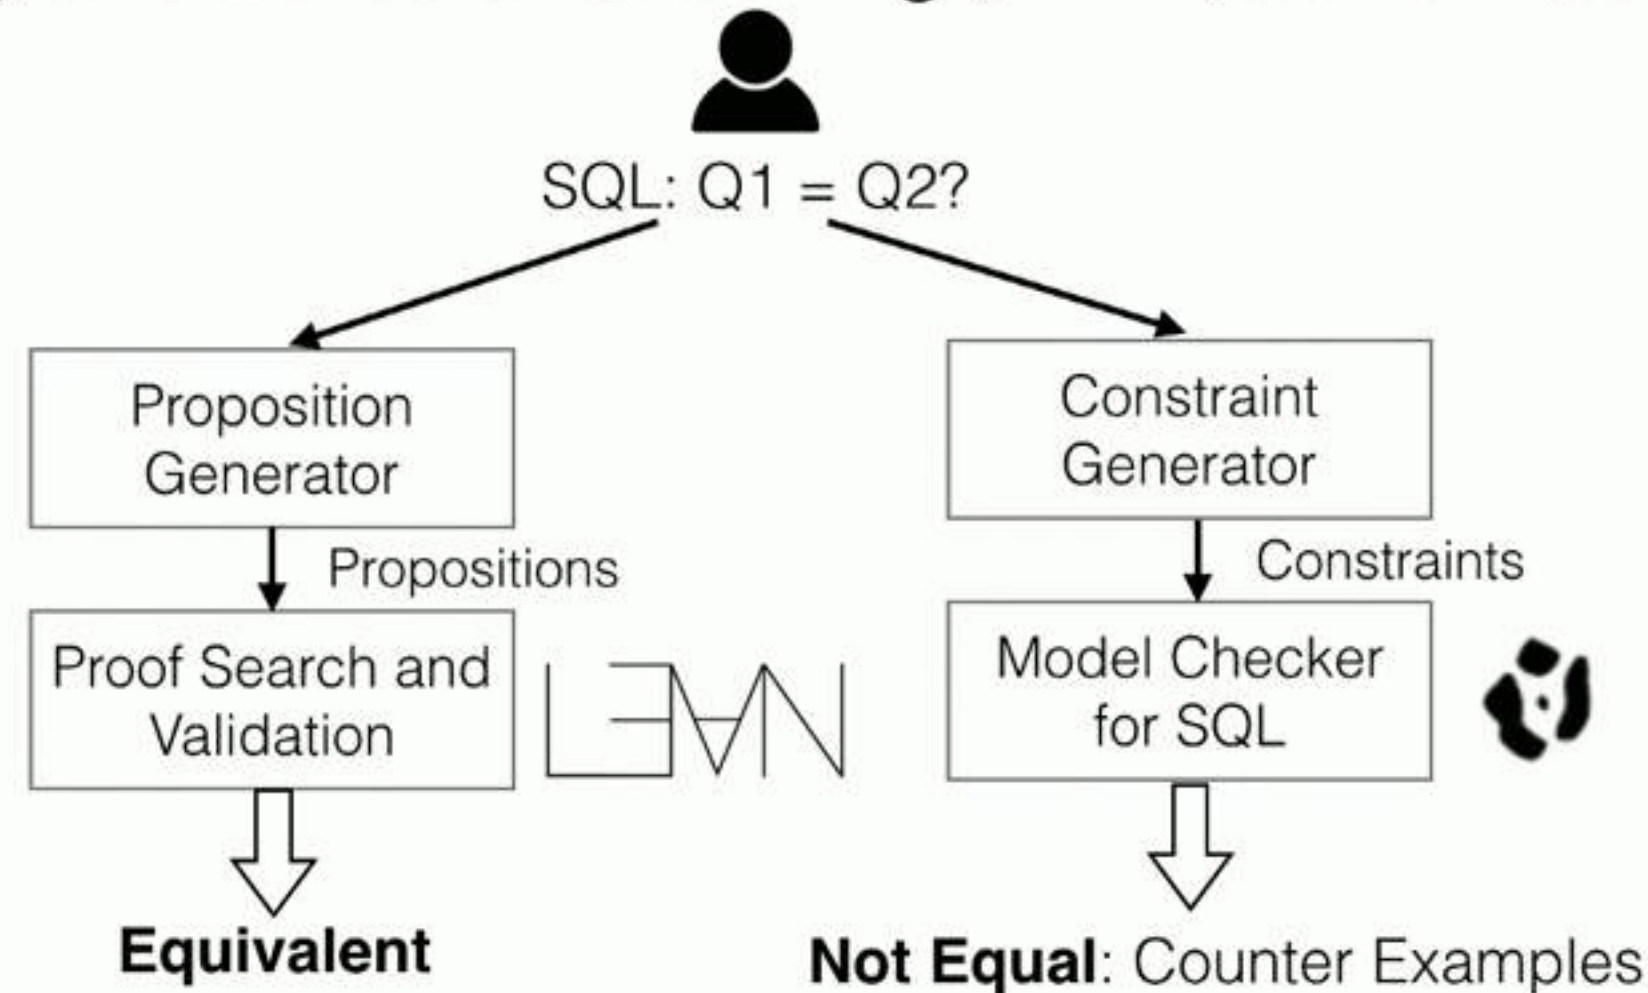
What about inequivalent queries?

Finding Counterexamples using Constraint Solver

[Chu et al., CIDR17]

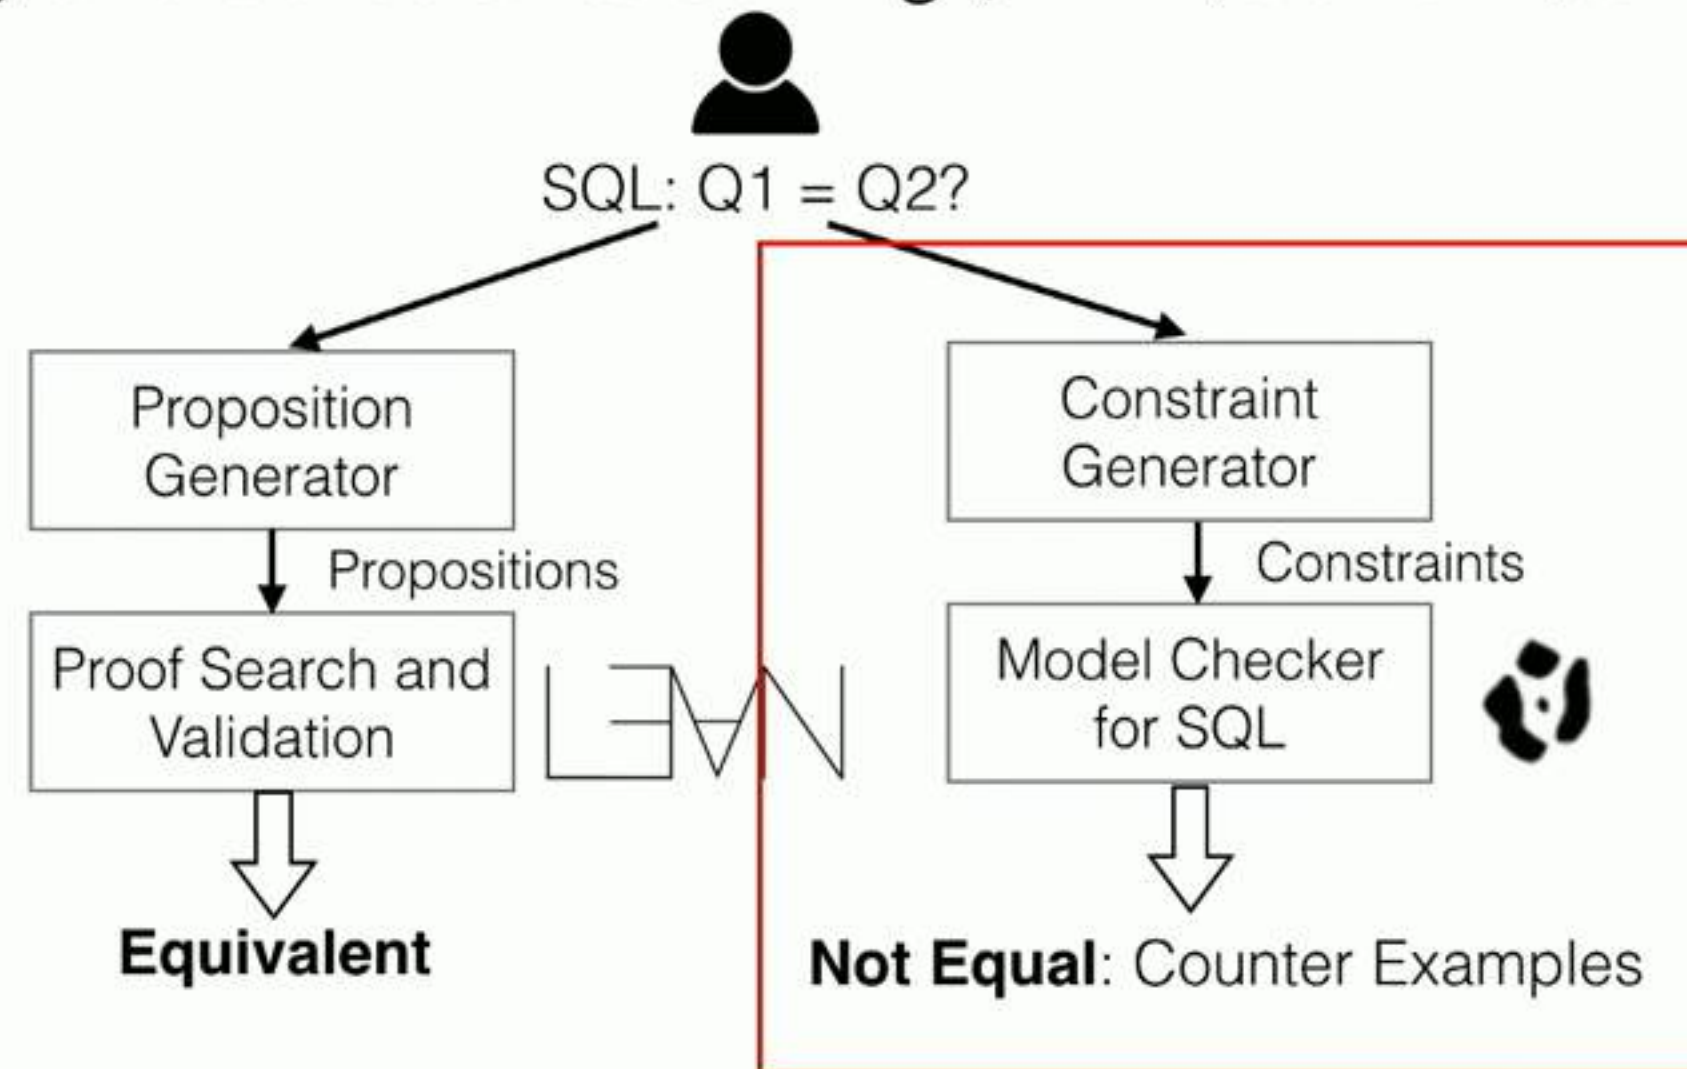
Cosette Architecture

- Cosette: An (almost) automated solver for SQL by *combining interactive theorem proving and constraints solving* (PLDI 17, SIGMOD17, CIDR17, VLDB18)

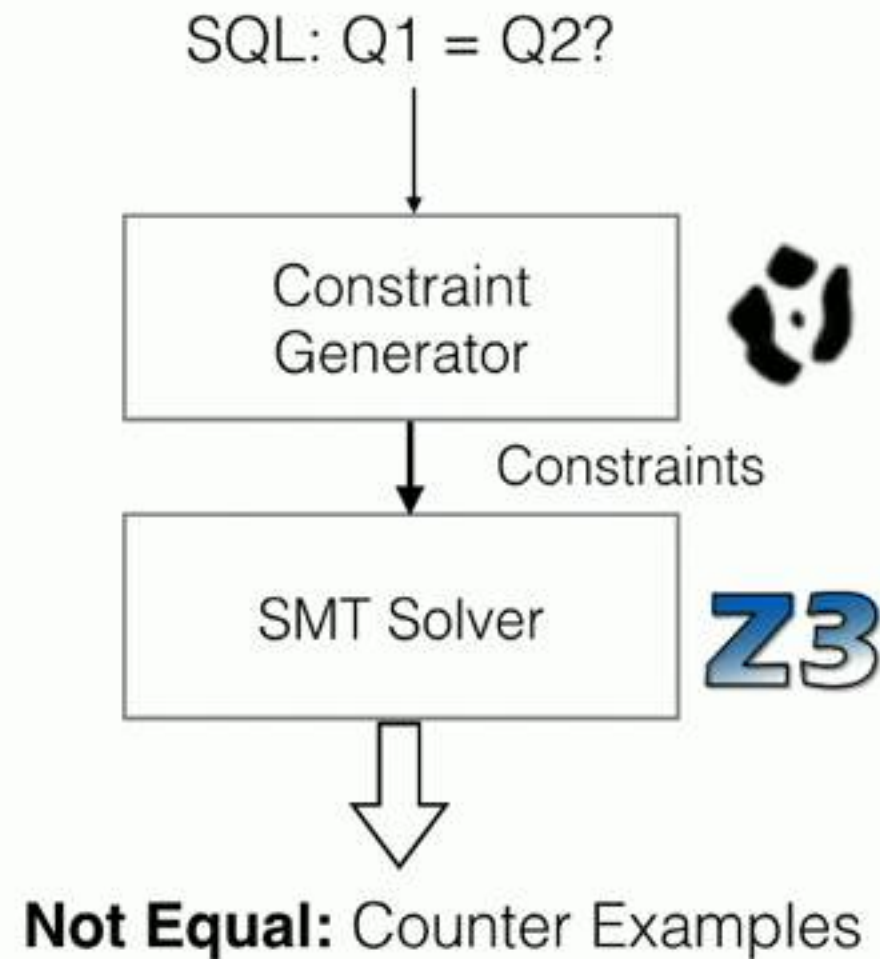


Cosette Architecture

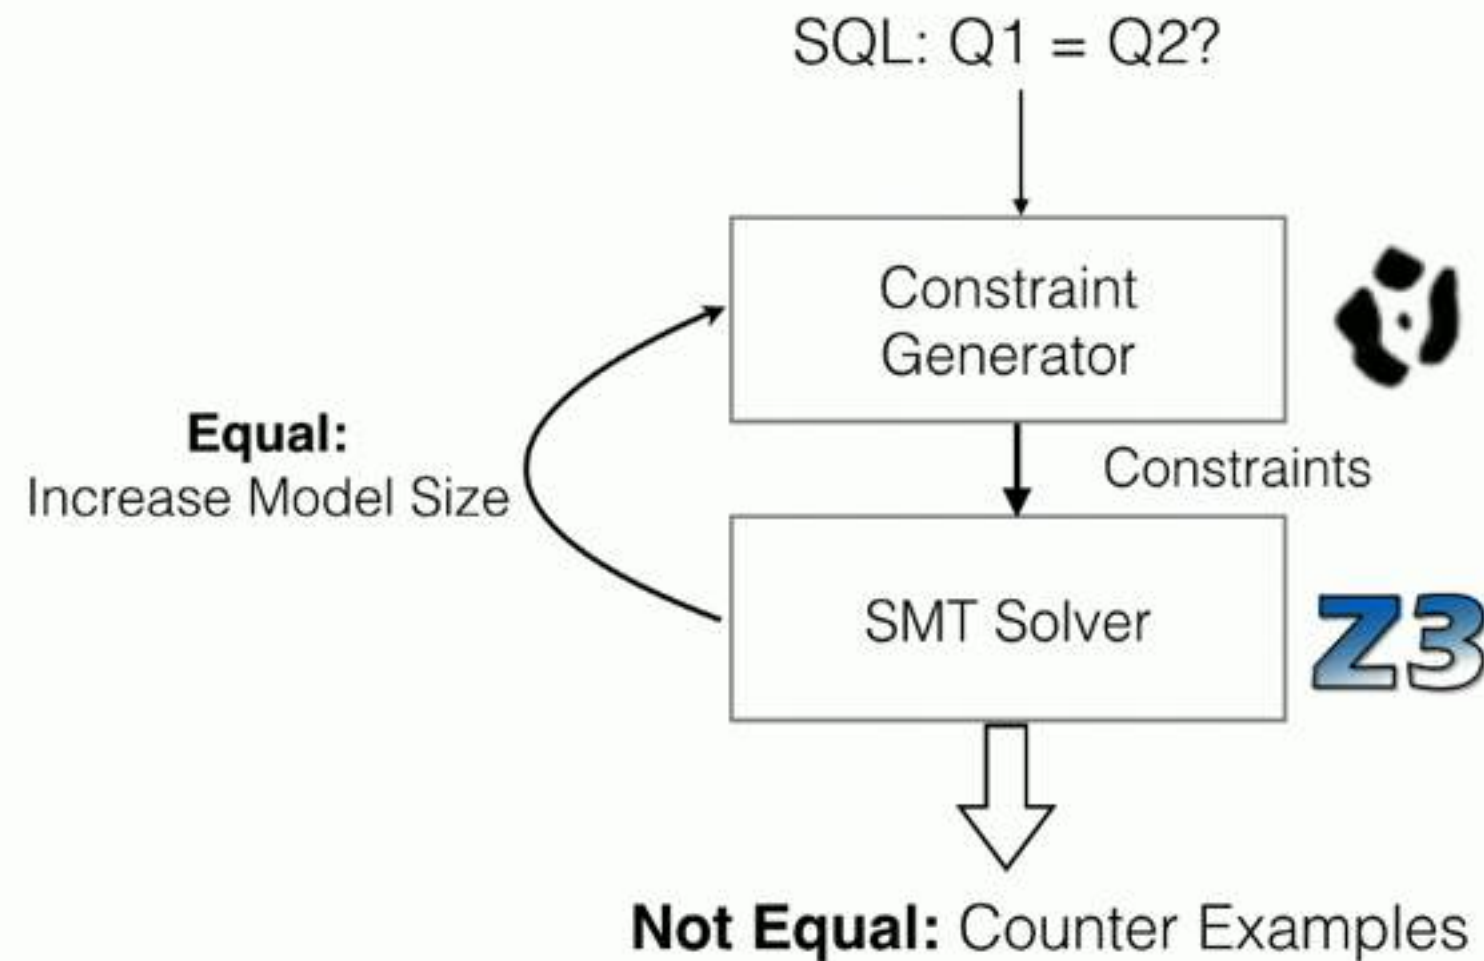
- Cosette: An (almost) automated solver for SQL by *combining interactive theorem proving and constraints solving* (PLDI 17, SIGMOD17, CIDR17, VLDB18)



Finding Counter Examples with SMT Solver



Finding Counter Examples with SMT Solver



Encoding SQL in Solvers

- A tuple as a list

`Tuple := List <Integer>`

Encoding SQL in Solvers

- A tuple as a list

`Tuple := List <Integer>`

- A relation as tuples tagged with multiplicity

`Relation := List <Pair<Tuple, Integer>>`

- A SQL query as constraints over symbolic values

Encoding SQL in Solvers

- A tuple as a list

SV
Tuple := List <Integer>

- A relation as tuples tagged with multiplicity

SV
Relation := List <Pair<Tuple, Integer>>

- A SQL query as constraints over symbolic values

Performance Optimizations

- Maximize concrete evaluation/ Minimize Symbolic Execution
- Symmetry breaking

Cosette

[Show API Key](#)

Cosette is an automated solver for checking equivalences of SQL queries. Check out [the Cosette Guide](#) on how to use Cosette. Your star and feedback are appreciated!

★ Star

444

! Issue

Example: Count Bug. ▼

```
1 schema s(pnum:int, shipdate:int);
2 schema p(pnum:int, qoh:int);
3
4 table parts(p);
5 table supply(s);
6
7 -- Kim, W. ACM Trans. Database System 1982
8 -- found incorrect 3 years later
9
10 query q1
11 `select x.pnum as xp
12   from parts x
13   where x.qoh = (select count(y.shipdate) as cnt
14                  from supply y
15                  where y.pnum = x.pnum AND y.shipdate < 10)`;
16
17 query q2
18 `select x.pnum as xp
19   from parts x, (select y.pnum as suppnum, count(y.shipdate) as
20                  ct
21                  from supply y where y.shipdate < 10
22                  group by y.pnum) temp
23   where x.qoh = temp.ct AND x.pnum = temp.suppnum`;
24 verify q1 q2;
```

Result

Submit

[Back to Cosette Website](#)

Cosette

[Show API Key](#)

Cosette is an automated solver for checking equivalences of SQL queries. Check out [the Cosette Guide](#) on how to use Cosette. Your star and feedback are appreciated!

[★ Star](#)

444

[! Issue](#)[Example: Count Bug. ▾](#)

```
1 schema s(pnum:int, shipdate:int);
2 schema p(pnum:int, qoh:int);
3
4 table parts(p);
5 table supply(s);
6
7 -- Kim, W. ACM Trans. Database System 1982
8 -- found incorrect 3 years later
9
10 query q1
11 `select x.pnum as xp
12 from parts x
```

Result

Organizations

ic Execution

Title & Bullets

[Change Master](#)

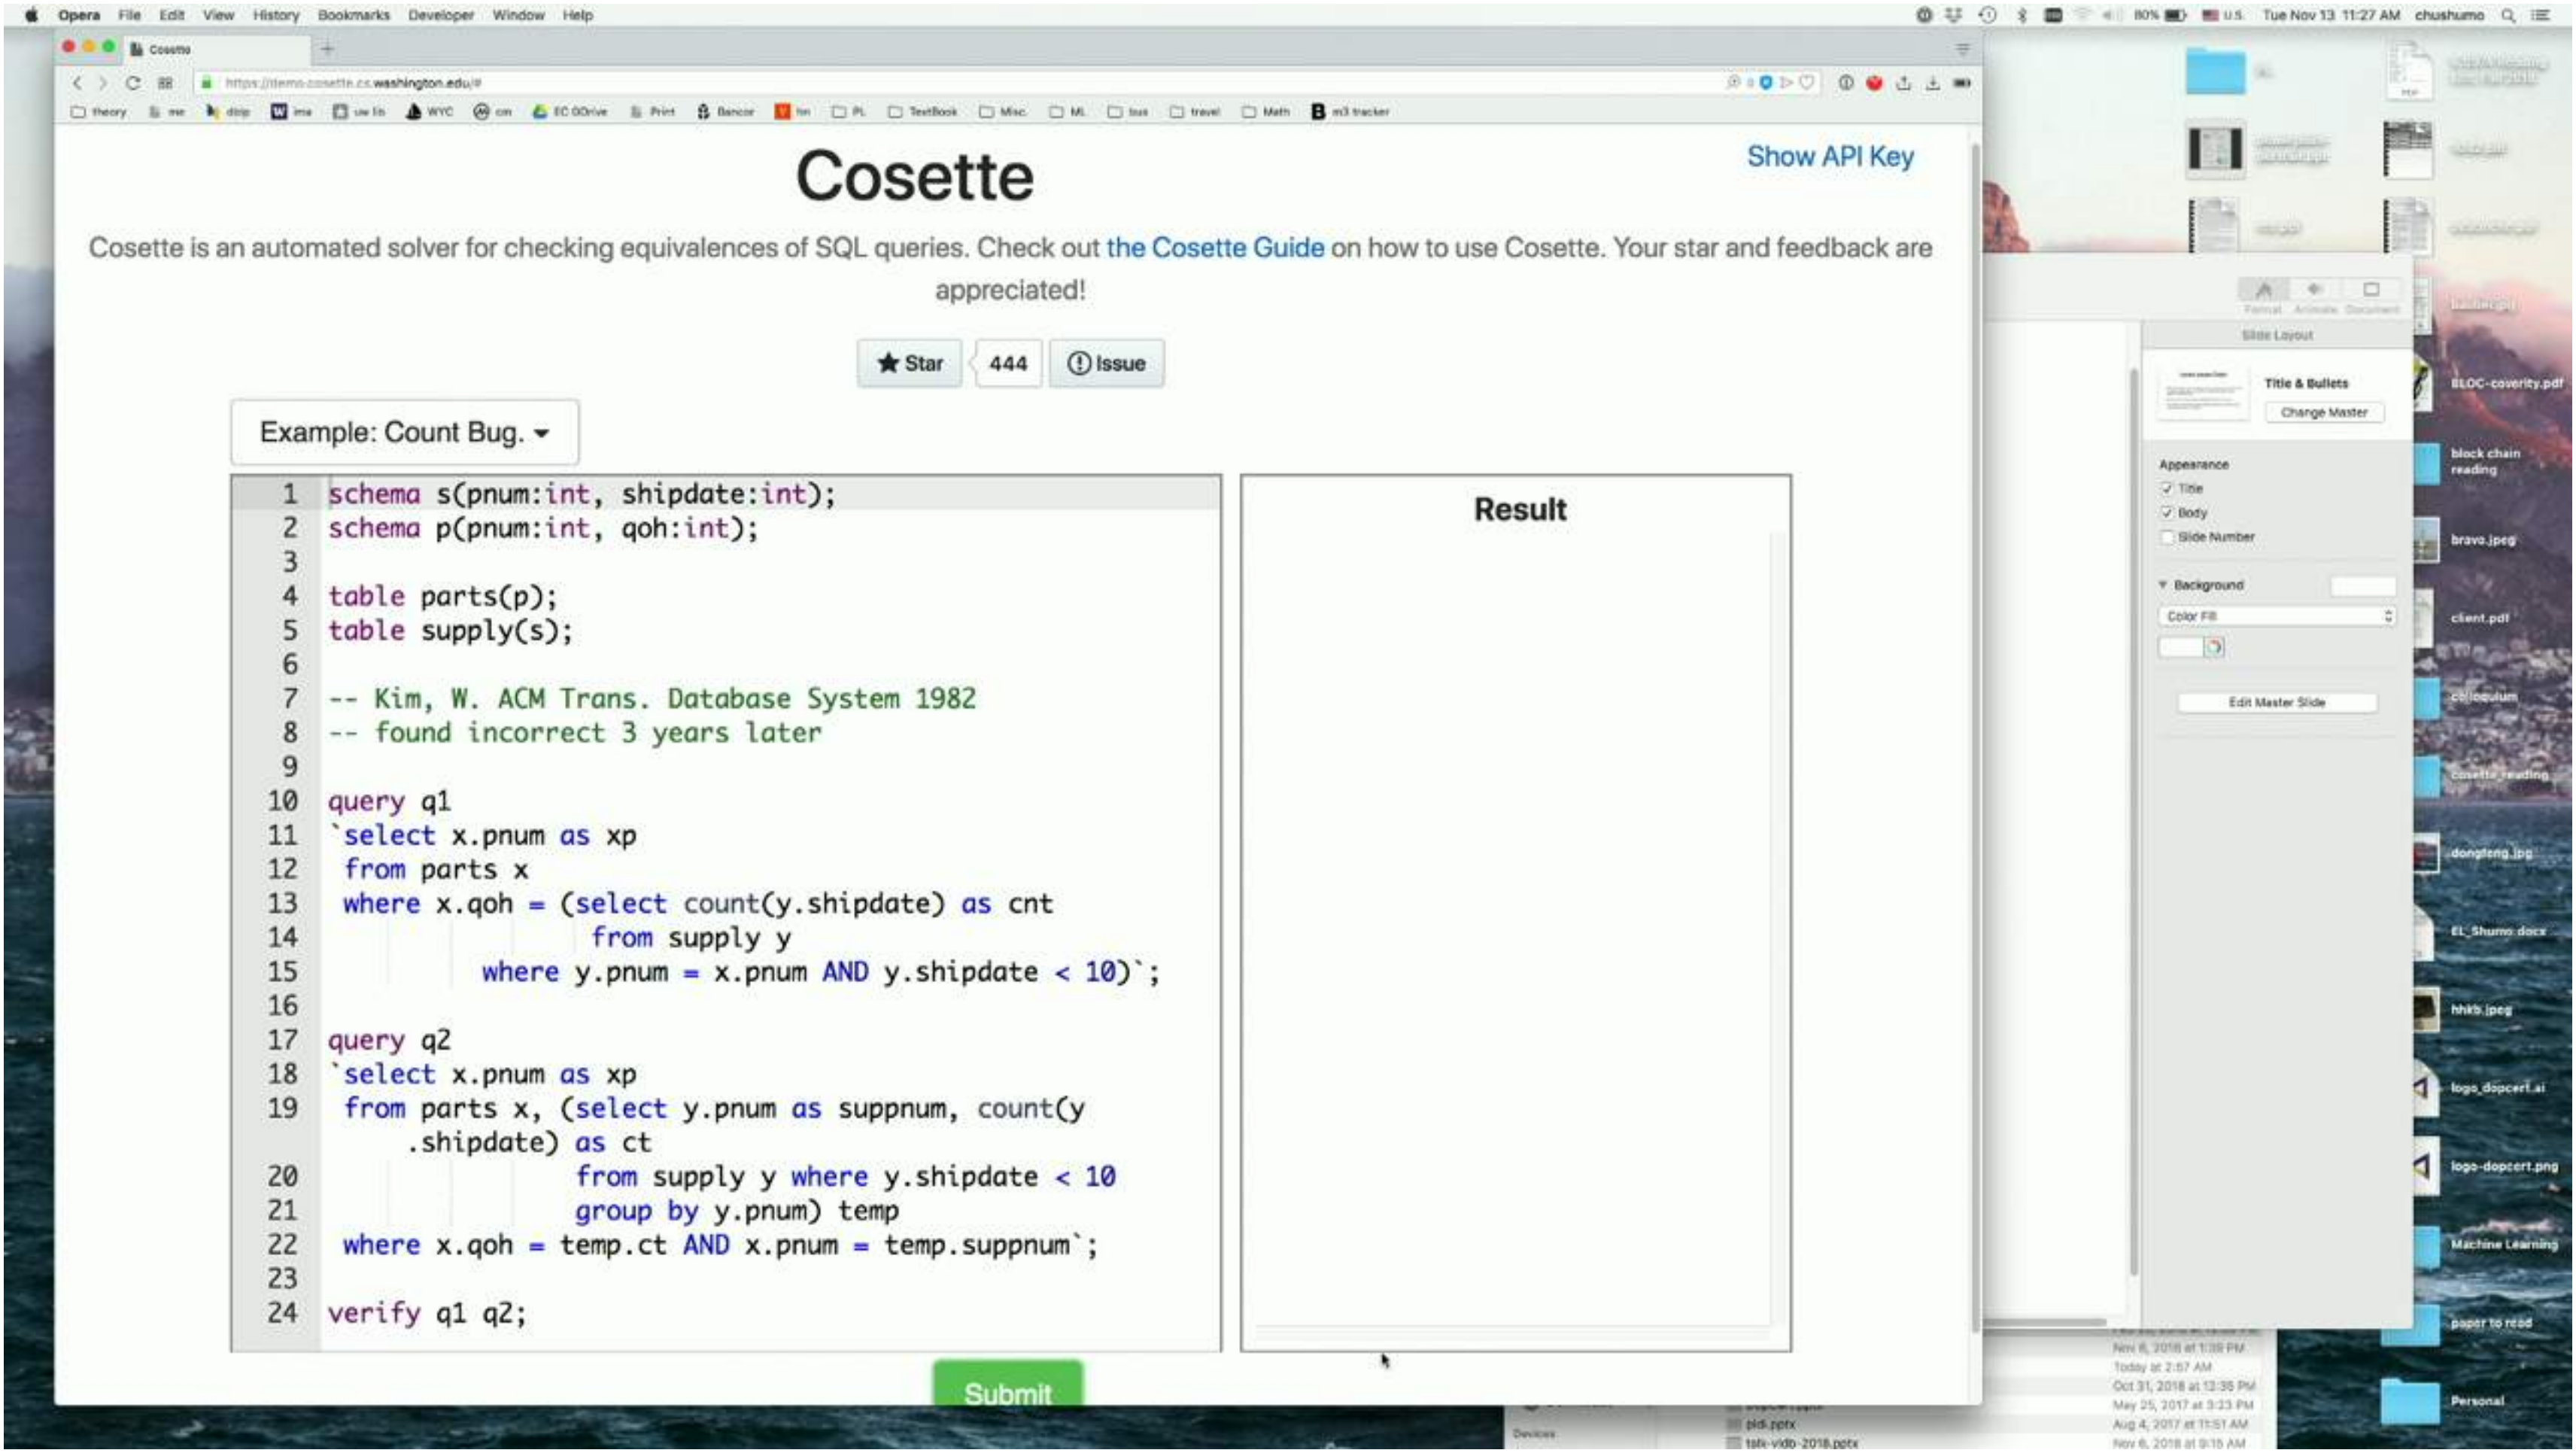
Appearance

- ☒ Title
- ☒ Body
- ☐ Slide Number

Background

Color Fill

[Edit Master Slide](#)



Cosette

[Show API Key](#)

Cosette is an automated solver for checking equivalences of SQL queries. Check out [the Cosette Guide](#) on how to use Cosette. Your star and feedback are appreciated!

[★ Star](#)

444

[! Issue](#)

Example: Count Bug. ▾

```
1 schema s(pnum:int, shipdate:int);
2 schema p(pnum:int, qoh:int);
3
4 table parts(p);
5 table supply(s);
6
7 -- Kim, W. ACM Trans. Database System 1982
8 -- found incorrect 3 years later
9
10 query q1
11 `select x.pnum as xp
12   from parts x
13   where x.qoh = (select count(y.shipdate) as cnt
14                  from supply y
15                  where y.pnum = x.pnum AND y.shipdate < 10)`;
16
17 query q2
18 `select x.pnum as xp
19   from parts x, (select y.pnum as suppnum, count(y
20                    .shipdate) as ct
21                  from supply y where y.shipdate < 10
22                  group by y.pnum) temp
23   where x.qoh = temp.ct AND x.pnum = temp.suppnum`;
24 verify q1 q2;
```

[Submit](#)

Result

Evaluating Cosette

- **Bug**: 3 real-world optimizer bugs
- **XData**: query and mutant pairs collected from XData, a test generation framework
- **Exams**: a set of questions from the undergraduate data management class
- **Rules**: 68 query rewrite rules from database literatures and real-world optimizers

Evaluating Cosette

- **Bug**: 3 real-world optimizer bugs
- **XData**: query and mutant pairs collected from XData, a test generation framework
- **Exams**: a set of questions from the undergraduate data management class
- **Rules**: 68 query rewrite rules from database literatures and real-world optimizers

Unequal
SQLs

Equivalent
SQLs

Evaluating Cosette

Dataset	Total #	Average time taken
Bugs	3	8.3s
XData	9	<1s
Exams	5	1.3s

Dataset	Equiv?	Total Number	Automatically Decided		Manually proved
			No.	Avg. Time	
Rules	Yes	68	62	< 20 s	6
Exams	Yes	4	3	< 1 s	1

Evaluating Cosette

Dataset	Total #	Average time taken
Bugs	3	8.3s
XData	9	<1s
Exams	5	1.3s

Rule No. 29:
400 LOC in
[Malecha et al., POPL 10]
to 15 LOC in Cosette

Dataset	Equiv?	Total Number	Automatically Decided		Manually proved
			No.	Avg. Time	
Rules	Yes	68	62	< 20 s	6
Exams	Yes	4	3	< 1 s	1

Cosette in Action

- Full stack: solver core, web service, online demo, automated grader
- Deployed for UW 344/544 automated grading since 2017
- SIGMOD 2017 Best Demo
- Top 1 Trending Racket Project in GitHub (July, 2017)

```
1  /* define schema s1,
2     here s1 can contain any number of attributes,
3     but it has to at least contain integer attributes
4     x and y */
5  schema s1(x:int, a:int, ??);
6
7  schema s2(a:int, ??);      -- define schema s2
8
9  table a(s1);               -- define table a using schema s1
10 table b(s2);               -- define table b using schema s1
11
12 query q1                   -- define query q1 on tables a
13     and b
14 `select distinct x.x as ax from a x, b y
15    where x.a = y.a`;
16
17 query q2                   -- define query q2 likewise
18 `select distinct x.x as ax from a x, a y, b z
19    where x.x = y.x and x.a = z.a`;
20
21 verify q1 q2;              -- does q1 equal to q2?
```

Conclusions and Takeaways

Cosette: The first practical SQL solver

- A new axiomatic semantics for SQL
- Semi-decision procedure for UCQ under bag/set with ICs
- Integrated interactive theorem proving and constraints solving techniques
- Automated reasoning brought by **Formal Methods + Domain Specific Semantics** *leads to more reliable, more optimized future data systems*
- Website: cosette.cs.washington.edu