

Interactive Debugging of Datalog Programs

André Pacak and Sebastian Erdweg



Sun 22 - Fri 27 October 2023 Cascais, Portugal

Attending ▾ Tracks ▾ Organization ▾ Search Series ▾


🏠 [SPLASH 2023 \(series\)](#) / [GPCE 2023 \(series\)](#) /

GPCE 2023

New Paper Category: Generative Pearls

...is an elegant essay about generative programming


Organizing Committee



Youyou Cong Publicity Chair
Tokyo Institute of Technology
Japan



Bernhard Rumpe General Chair
RWTH Aachen University
Germany



Amir Shaikhha Program Chair
University of Edinburgh
United Kingdom

Important Dates AoE (UTC-12h)

Mon 3 Jul 2023 new

Abstract submission

Fri 7 Jul 2023 new

Paper submission

Wed 23 Aug 2023 new

Interactive Debugging of Datalog Programs


André Pacak and Sebastian Erdweg

Datalog 101

Head Body

←←

$\text{path}(X, Y) \text{ :- edge}(X, Y) \text{ .}$
 $\text{path}(X, Y) \text{ :- edge}(X, Z), \text{ path}(Z, Y) \text{ .}$



IDB = derived facts

$\text{path}(1, 2) \text{ .}$ $\text{path}(1, 3) \text{ .}$ $\text{path}(1, 4) \text{ .}$ $\text{path}(1, 1) \text{ .}$
 $\text{path}(2, 3) \text{ .}$ $\text{path}(2, 4) \text{ .}$ $\text{path}(2, 1) \text{ .}$ $\text{path}(2, 2) \text{ .}$
 $\text{path}(3, 4) \text{ .}$ $\text{path}(3, 1) \text{ .}$ $\text{path}(3, 2) \text{ .}$ $\text{path}(3, 3) \text{ .}$

EDB = given facts

←←

$\text{edge}(1, 2) \text{ .}$
 $\text{edge}(2, 3) \text{ .}$
 $\text{edge}(3, 4) \text{ .}$
 $\text{edge}(3, 1) \text{ .}$

Datalog gem 1: Computable least fixpoints

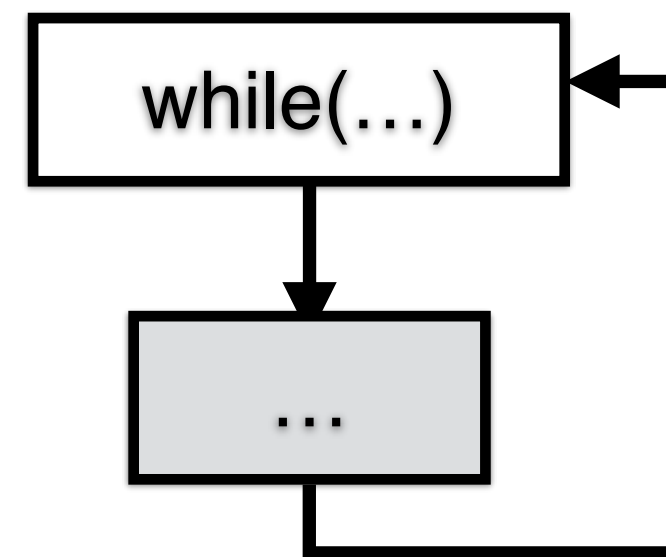
// transitive closure

```
reachable(Prog, From, To) :- flow(Prog, From, To).
```

```
reachable(Prog, From, To) :- flow(Prog, From, Inter), reachable(Prog, Inter, To).
```

// loop detection

```
looping(Stm) :- reachable(Prog, Stm, Stm).
```



// data flow analysis

```
entry(Prog, Stm, X, V) :- flow(Prog, Pred, Stm), join(Pred, X, V).
```

```
exit(Prog, Stm, X, V) :- assign(Stm, X, E), eval(E, V).
```

```
exit(Prog, Stm, X, V) :- not assign(Stm, _, _), entry(Prog, Stm, X, V)
```

Datalog gem 2: Efficient incremental computing

IncA: Incremental static analysis framework

- Analyses implemented in Datalog
- Code changes fed to incremental Datalog engine

Language	Analysis	Programs	Non-inc. time	Inc. time	Speedup
Java	FindBugs	mbeddr importer (10k LoC)	n/a	7 ms	n/a
C	flow-sensitive points-to	Toyota ITC code ² (15k LoC)	5800 ms	23.3 ms	249x
C	well-formedness checks	Smart Meter (44k LoC)	209 ms	12.8 ms	16.3x
Jimple	Strong-update points-to	[GTruth (9k), Gson (14k), PGSQL]	6500–64300 ms	1–10 ms	650–6430 x
Jimple	String analysis	[JDBC (45k), BerkleyDB (70k)]	13500–20400 ms	2 ms	6500–10000x
Jimple	Constant propagation	[antlr (22k), emma (26k)]	5000–23000 ms	1–3 ms	1600–7600x
Jimple	Interval analysis	[pmd (61k), ant (105k)]	3000–23000 ms	1–6 ms	500–3800x

Increasing Datalog Complexity => Need Better Dev Tools

Answer questions from users and developers:

Why does variable x point to heap location a?

?- varpointsto("x",a)

Why does y not point to a?

?- varpointsto("y",a)

What locations does z point to?

?- varpointsto("z",A)

Are there alternative derivations?


?- varpointsto("x",a)

```
var x = new Object // loc a
var y = new Object // loc b
var z = new Object // loc c

if (x == null) {
    y = x
}
if (Random.nextBool()) {
    z = x
} else {
    z = y
}
x = z
```

Bottom-up semantics

$\text{path}(X, Y) \text{ :- edge}(X, Y).$
 $\text{path}(X, Y) \text{ :- edge}(X, Z), \text{path}(Z, Y).$



iter = 0

$\text{edge}(1, 2).$
 $\text{edge}(2, 3).$
 $\text{edge}(3, 4).$
 $\text{edge}(3, 1).$

iter = 1

$\text{path}(1, 2).$
 $\text{path}(2, 3).$
 $\text{path}(3, 4).$
 $\text{path}(3, 1).$

iter = 2

$\text{path}(1, 3).$
 $\text{path}(2, 4).$
 $\text{path}(2, 1).$
 $\text{path}(3, 2).$

iter = 3

$\text{path}(1, 4).$
 $\text{path}(1, 1).$
 $\text{path}(2, 2).$
 $\text{path}(3, 3).$

BOTTOM-UP BEATS TOP-DOWN FOR DATALOG

Jeffrey D. Ullman
 Stanford University


ABSTRACT

We show that for any safe datalog program \mathcal{P}_1 and any query Q (predicate of \mathcal{P}_1 with some bound arguments), there is another safe datalog program \mathcal{P}_2 that produces the answer to Q and takes no more time when evaluated

2. Safety of rules, that is, every variable in the head appears in some ordinary (not a built-in arithmetic comparison) subgoal.
3. Semi-naive evaluation, that is, differential computation of the least fixed point.

Bottom-up semantics

$\text{path}(X, Y) \text{ :- edge}(X, Y).$
 $\text{path}(X, Y) \text{ :- edge}(X, Z), \text{path}(Z, Y).$



iter = 0

$\text{edge}(1, 2).$
 $\text{edge}(2, 3).$
 $\text{edge}(3, 4).$
 $\text{edge}(3, 1).$

iter = 1

$\text{path}(1, 2).$
 $\text{path}(2, 3).$
 $\text{path}(3, 4).$
 $\text{path}(3, 1).$

iter = 2

$\text{path}(1, 3).$
 $\text{path}(2, 4).$
 $\text{path}(2, 1).$
 $\text{path}(3, 2).$

iter = 3

$\text{path}(1, 4).$
 $\text{path}(1, 1).$
 $\text{path}(2, 2).$
 $\text{path}(3, 3).$

Why is there a path from 3 to 3

?- $\text{path}(3, 3)$

Why is there no $\text{path}(4, 1)$

?- $\text{path}(4, 1)$

What paths start at 3 and why

?- $\text{path}(3, Y)$

Are there alternative derivations for


?- $\text{path}(3, 4)$

Provenance-based debugging

$$\begin{array}{c}
 \text{R2} \frac{\text{edge}(3, 1)}{\text{path}(3, 3)} \\
 \text{R2} \frac{\text{edge}(1, 2) \quad \text{R1} \frac{\text{edge}(2, 3)}{\text{path}(2, 3)}}{\text{path}(1, 3)}
 \end{array}$$


Bottom-up semantics — forward chaining

- Starts with the given facts
- Tries to derive new facts by applying rules
- Continues until no new facts can be found




Top-down semantics — backward chaining

- Starts with a query and operates demand-driven
- Tries to answer that query by applying rules
- Continues until query is fully answered
- Yields same result as bottom-up semantics
- Very similar to standard programming experience



Top-down stepping = Interactive debugging

Why is there a path from 3 to 3



$[X=3, Y=3]$ $[X=3, Y=3]$
path(X, Y) :- edge(X, Y) . ~~$[X=3, Y=3]$~~ $\{ [X=3, Y=3, Z=4],$
 $[X=3, Y=3, Z=1] \}$

$[X=3, Y=3]$ $[X=3, Y=3]$ $[X=3, Y=3, Z=1]$
path(X, Y) :- edge(X, Z), path(Z, Y) . $[X=3, Y=3, Z=1]$ ✓

$\{ [X=4, Y=3],$
 $[X=1, Y=3] \}$
path(X, Y) :- edge(X, Y) . $\{ [X=4, Y=3],$ ~~$[X=1, Y=3]$~~ $\}$

$\{ [X=4, Y=3],$ $\{ [X=4, Y=3],$
 $[X=1, Y=3] \}$ $[X=1, Y=3] \}$ $\{ [X=1, Y=3, Z=2] \}$

path(X, Y) :- edge(X, Z), path(Z, Y) . $[X=1, Y=3, Z=2]$ ✓

Top-down stepping = Interactive debugging

Problem 1: No top-down SOS exists for Datalog

- Existing algorithms: Query-Subquery (QSQ) Iterative and Recursive
- Beware: Lots of literature presents incomplete QSQ algorithms

Problem 2: Top-down performance is poor

- Semi-naïve evaluation and n-ary joins beat top-down
- (Virtually) all serious Datalog implementations use bottom-up

BOTTOM-UP BEATS TOP-DOWN FOR DATALOG

Jeffrey D. Ullman
Stanford University

ABSTRACT

We show that for any safe datalog program \mathcal{P}_1 and any query Q (predicate of \mathcal{P}_1 with some bound arguments), there is another safe datalog program \mathcal{P}_2 that produces

2. Safety of rules, that is, every variable in the head appears in some ordinary (not a built-in arithmetic comparison) subgoal.
3. Semi-naive evaluation, that is, differential compu-

Datalog syntax

(Datalog programs)	$D ::= \bar{r}, \bar{f}$
(rules)	$r ::= p(\bar{X}) :- \bar{a}.$
(atoms)	$a ::= p^s(\bar{t}) \mid \text{edb } p^s(\bar{t}) \mid t = t$
(signs)	$s ::= + \mid -$
(terms)	$t ::= c \mid X$
(facts)	$f ::= p(\bar{c}).$
(constants)	c
(variables)	X

(queries) $Q ::= \mathbf{sq}(p^s(\bar{t}), v, v, v, r \vee \dots \vee r) \mid v^s$

- original predicate call
- query arguments
- partial subquery results
- bindings of current rule
- remaining rules

Reduction relations

(rule reduction)	$v \vdash r \xrightarrow{R} r \vdash v$
(atom reduction)	$v \vdash a \xrightarrow{A} Q$
(query reduction)	$Q \xrightarrow{Q} Q$

Top-down Datalog SOS

(rule reduction)

$$v \vdash r \longrightarrow^R r \dashv v$$

$$\text{R-Step} \frac{v_{sup} \vdash a \longrightarrow^A a'}{v_{sup} \vdash p(\bar{X}) :- a, \bar{a}s. \longrightarrow^R p(\bar{X}) :- a', \bar{a}s. \dashv v_{sup}}$$

$$\text{R-Merge} \frac{}{v_{sup} \vdash p(\bar{X}) :- v^s, \bar{a}s. \longrightarrow^R p(\bar{X}) :- \bar{a}s. \dashv \text{merge}(v_{sup}, v^s)}$$

$$\text{R-Result} \frac{}{v_{sup} \vdash p(\bar{X}) :- \epsilon. \longrightarrow^R \Pi_{\bar{X}}(v_{sup}) \dashv v_{sup}}$$

Top-down Datalog SOS

(atom reduction)

$$v \vdash a \rightarrow^A Q$$

$$\text{A-Into} \frac{v_a = \text{eval}(\text{cols}(p), \bar{t}, v_{sup}) \quad (\Gamma', v'_a) = \text{pushQuery}(\Gamma, p, v_a) \quad v'_a \text{ not empty}}{v_{sup} \vdash p^s(\bar{t}) \rightarrow^A \mathbf{sq}(p^s(\bar{t}), v'_a, \text{table}(\text{cols}(p), \emptyset), v'_a, \text{rules}(p))} \Gamma := \Gamma'$$

Top-down Datalog SOS

(query reduction)

$$Q \rightarrow^Q Q$$

$$\text{Q-Iterate} \frac{v_r \not\subseteq IDB(p) \quad IDB' = IDB \uplus p \mapsto (v_r \cup IDB(p))}{\mathbf{sq}(p^s(\bar{t}), v_q, v_r, v_{sup}, \epsilon) \rightarrow^Q \mathbf{sq}(p^s(\bar{t}), v_q, \text{table}(\text{cols}(p), \emptyset), v_q, \text{rules}(p))} IDB := IDB'$$

$$\text{Q-Stable} \frac{v_r \subseteq IDB(p) \quad (\Gamma', v_a) = \text{popQuery}(\Gamma, p, v_q)}{\mathbf{sq}(p^s(\bar{t}), v_q, v_r, v_{sup}, \epsilon) \rightarrow^Q \rho_{\bar{t}/\text{cols}(p)}(IDB(p) \bowtie v_a)^s} \Gamma := \Gamma'$$

Implementation and Performance

Implemented in InCA


Evaluated on Doop

?- methLookup(C, "accept")

?- supertypes({"C","D"}, S)

?- varPointsTo("x", a)


Time per step-into interaction per rule




Performance problem

Simulate step-over and resume through repeated step-into?

Evaluated on synthetic cyclic graphs, computing path relation




Running time for different graph sizes



Hybrid semantics

$$\text{A-Over} \frac{v_a = \text{eval}(\text{cols}(p), \bar{t}, v_{sup})}{v_{sup} \vdash p^s(\bar{t}) \rightarrow^A \rho_{\bar{t}/\text{cols}(p)}(\text{BU}(p) \bowtie v_a)^s}$$

Top-down stepping = Interactive debugging



[X=1] path(X, Y) :- [X=1] edge(X, Y). [X=1, Y=2] ✓

[X=1] path(X, Y) :- [X=1] edge(X, Z), [X=1, Z=2] path(Z, Y).
 { [X=1, Y=1, Z=2],
 [X=1, Y=2, Z=2],
 [X=1, Y=3, Z=2],
 [X=1, Y=4, Z=2],
 [X=1, Y=5, Z=2] }

[X=2] path(X, Y) $\xrightarrow{\text{over}}$ { [X=2, Y=1],
 [X=2, Y=2],
 [X=2, Y=3],
 [X=2, Y=4],
 [X=2, Y=5] }

Hybrid semantics

$$\text{A-Over} \frac{v_a = eval(cols(p), \bar{t}, v_{sup})}{v_{sup} \vdash p^s(\bar{t}) \rightarrow^A \rho_{\bar{t}/cols(p)}(BU(p) \bowtie v_a)^s}$$


$$\text{A-OverRecursive} \frac{recursiveCall(p^s(\bar{t})) \quad v_a = eval(cols(p), \bar{t}, v_{sup})}{v_{sup} \vdash p(\bar{t})^s \rightarrow^A \rho_{\bar{t}/cols(p)}(BU(p, \Gamma) \bowtie v_a \cup IDB(p) \bowtie v_a)^s}$$

$BU(p, \Gamma)$ = incrementally maintained, blacklist-aware bottom-up database


Performance problem

Simulate step-over and resume through repeated step-into?

Evaluated on synthetic cyclic graphs, computing path relation



Running time for different graph sizes



Step-over performance


Implemented in InCA

Evaluated on Doop varPointsTo

Step over, except:

- a) VarPointsTo
- b) VarPointsTo, StaticFieldPointsTo
- c) VarPointsTo, InstanceFieldPointsTo
- d) VarPointsTo, Reachable

Time per step-over interaction



Interactive Debugging of Datalog

Top-down semantics: program exploration

Hybrid semantics: efficient step-over/resume

Vision: Datalog as IR with rich tooling

