## **Datalog Revival**

#### Serge Abiteboul INRIA Saclay & ENS Cachan





## Datalog history

Started in 77: logic and database workshop

Simple idea: add recursion to positive FO queries

Blooming in the 80<sup>th</sup>

Logic programming was hot

Industry was not interested:

 "No practical applications of recursive query theory ... have been found to date." Hellerstein and Stonebraker (Readings in DB Systems)

Quasi dead except local resistance [e.g., A., Gottlob]

Revival in this century



## Organization

- Datalog
- Datalog evaluation
- Datalog with negation
- Datalog revival
- Conclusion

# Datalog

#### Limitation of relational calculus

G a graph: G(0,1), G(1,2), G(2,3), ... G(10,11) Is there a path from *0* to *11* in the graph?



k-path  $\exists x_1 \dots x_k$  (G(0,x\_1) $\land$ G(x<sub>1</sub>,x<sub>2</sub>) $\land$ ... $\land$ G(x<sub>k-1</sub>,x<sub>k</sub>)  $\land$ G(x<sub>k</sub>,11)) Path of unbounded length: **infinite** formula



G(2,3) $T(x, y) \leftarrow G(x, z), T(z, y)$ 

fact rule

datalog rule : 
$$R_1(u_1) \leftarrow R_2(u_2), \ldots, R_n(u_n)$$
 for  $n \ge 1$   
head body

- Each u<sub>i</sub> is a vector of terms
- Safe: each variable occurring in head must occur in body
- Intentional relation: occurs in the head
- Extensional relation: does not

#### Datalog program

- 1. G(0,1), G(1,2), G(2,3), ... G(10,11)
- 2.  $T(x, y) \leftarrow G(x, y)$
- 3.  $T(x, y) \leftarrow G(x, z), T(z, y)$
- 4. Ok()  $\leftarrow$  T(0, 11)

edb(P) = {G}
idb(P) = {T,Ok}
program P

#### Datalog program

- 1. G(0,1), G(1,2), G(2,3), ... G(10,11)
- 2.  $T(x, y) \leftarrow G(x, y)$   $T(10, 11) \leftarrow G(10, 11)$
- 3.  $T(x, y) \leftarrow G(x, z), T(z, y)$
- 4. Ok()  $\leftarrow$  T(0, 11)

 $T(\theta, 11) \leftarrow G(\theta, 10) T(101) 1$ Ok()  $\leftarrow T(0, 11)$ 

Rule 2: v(x)=10 & v(y) = 11 $rac{10,11}{2}$ Rule 3: v(x)=9, v(z)=10 & v(y)=11 $rac{10,11}{2}$ 

Rule 3: v(x)=0, v(z)=1 & v(y)=11Rule 4: v(x)=0, v(y)=11

☞ T(0,11) ☞ <mark>Ok(</mark>)

...

#### **Model semantics**

View P as a first-order sentence  $\Sigma_{P}$  describing the answer

- Associate a formula to each rule  $R_{1}(u_{1}) \leftarrow R_{2}(u_{2}), \dots, R_{n}(u_{n}) :$   $\forall x_{1}, \dots, x_{m}(R_{2}(u_{2}) \land \dots \land R_{n}(u_{n}) \Longrightarrow R_{1}(u_{1}))$ where  $x_{1}, \dots, x_{m}$  are the variables occurring in the rule  $P = \{r_{1}, \dots, r_{n}\}, \Sigma_{P} = r_{1} \land \dots \land r_{n}$ 

The **semantics** of **P** for a database **I**, denoted **P(I)**, is the **minimum model of**  $\Sigma_{P}$  **containing I** 

Does it always exist? How can it be computed?

#### **Example:** Transitive closure

```
G(0,1), G(1,2), G(2,3)
T(x,y) \leftarrow G(x,y)
T(x,y) \leftarrow G(x,z), T(z,y)
```



## Existence of P(I)

There exists at least one such model: the largest instance one can build with the constants occurring in I and P is a model of P that includes I – B(I,P)

P(I) always exists: it is the intersection of all models of P that include I over the constants occurring in I and P

How can it be computed?

#### **Fixpoint semantics**

A fact A is an *immediate consequence* for K and P if

- 1. A is an extensional fact in K, or
- 2. for some instantiation  $A \leftarrow A_1, \ldots, A_n$  of a rule in P, each  $A_i$  is in K

Immediate consequence operator:

T<sub>P</sub>(K) = { immediate consequences for K and P }

Note: T<sub>P</sub> is monotone

Fixpoint semantics – continued

P(I) is a fixpoint of  $T_P - That$  is:  $T_P(P(I)) \subseteq P(I)$ Indeed, P(I) is the least fixpoint of  $T_P$  containing I

Yields a means of computing P(I)

 $I \subseteq T_{p}(I) \subseteq T_{p}^{2}(I) \subseteq \dots \subseteq T_{p}^{i}(I) = T_{p}^{i+1}(I) = P(I) \subseteq B(I,P)$ 

## **Proof theory**

- Proof technique: SLD resolution
- A fact A is in P(I) iff there exists a proof of A

#### Static analysis

#### Hard

- Deciding containment ( $P \subseteq P'$ ) is undecidable
- Deciding equivalence is undecidable
- Deciding boundedness is undecidable
  - There exists k such that for any I, the fixpoint converges in less than k stages
- So, optimization is hard

# Datalog evaluation by example

# More complicated example: Reverse same generation

| up                                                    |   | flat |   | down |   |  |  |
|-------------------------------------------------------|---|------|---|------|---|--|--|
| а                                                     | е | g    | f | Ι    | f |  |  |
| а                                                     | f | m    | n | m    | f |  |  |
| f                                                     | m | m    | 0 | g    | b |  |  |
| g                                                     | n | р    | m | h    | С |  |  |
| h                                                     | n |      |   | i    | d |  |  |
| i                                                     | 0 |      |   | р    | k |  |  |
| j                                                     | 0 |      |   |      |   |  |  |
| $sg(x,y) \leftarrow flat(x,y)$                        |   |      |   |      |   |  |  |
| $sg(x,y) \leftarrow up(x,x1), rsg(y1,x1), down(y1,y)$ |   |      |   |      |   |  |  |





4/29/2013

18

#### Naive algorithm

**Fixpoint**  $rsg_0 = \emptyset$  $rsg_{i+1} = flat \cup rsg_i \cup \pi_{16}(\sigma_{2=4}(\sigma_{3=5}(up \times rsg_i \times down)))$ Program rsg :=  $\emptyset$ ; repeat rsg := flat  $\cup$  rsg  $\cup \pi_{16}(\sigma_{2=4}(\sigma_{3=5}(up \times rsg \times down)))$ until fixpoint

#### Semi-naive

 $\begin{array}{ll} \Delta_1(\mathbf{x},\mathbf{y}) &\leftarrow \operatorname{flat}(\mathbf{x},\mathbf{y}) \\ \Delta_{i+1}(\mathbf{x},\mathbf{y}) \leftarrow up(\mathbf{x},\mathbf{x1}), \Delta_i(\mathbf{y1},\mathbf{x1}), \operatorname{down}(\mathbf{y1},\mathbf{y}) \\ \operatorname{Compute} U\Delta_1 \end{array}$ 

Program

- Converges to the answer
- Not recursive & not a datalog program
- Still redundant to avoid it:

 $\Delta_{i+1}(x, y) \leftarrow up(x, x1), \Delta_i(y1, x1), down(y1, y), \neg \Delta_i(x, y)$ 





#### Semi-naïve (end)

More complicated if the rules are not linear  $T(x, y) \leftarrow G(x, y)$  $T(x, y) \leftarrow T(x, z), T(z, y)$ 

- $\Delta_1(x, y) \leftarrow G(x, y)$
- anc1 :=  $\Delta 1$
- tempi+1(x, y)  $\leftarrow \Delta i(x, z)$ , anci(z, y)
- tempi+1(x, y)  $\leftarrow$  anci(x, z),  $\Delta i(z, y)$
- $\Delta^{i+1} := temp^{i+1} anc^i$
- anci+1 := anci  $\cup \Delta^{i+1}$

## And beyond

Start from a program and a query

 $rsg(x,y) \leftarrow flat(x,y)$ 

 $rsg(x,y) \leftarrow up(x,x1), rsg(y1,x1), down(y1,y)$ 

query(y)  $\leftarrow$  rsg(a, y)

**Optimize to avoid deriving useless facts** 

Two competing techniques that are roughly equivalent

- Query-Sub-Query
- Magic Sets

#### Magic Set

```
rsg^{bf}(x, y) \leftarrow input_{rsg^{bf}}(x), flat(x, y)
     rsgfb(x, y) \leftarrow input_{rsg}^{fb}(y), flat(x, y)
    sup31(x, x1) \leftarrow input_rsg^{bf}(x), up(x, x1)
    sup32(x, y1) \leftarrow sup31(x, x1), rsg^{fb}(y1, x1)
    rsg^{bf}(x, y) \leftarrow sup 32(x, y1), down(y1, y)
    sup41(y, y1) \leftarrow input_rsg^{fb}(y), down(y1, y)
    sup42(y, x1) \leftarrow sup41(y, y1), rsg^{bf}(y1, x1)
     rsg^{fb}(x, y) \leftarrow sup42(y, x1), up(x, x1)
    input_rsg<sup>bf</sup>(x1) \leftarrow sup31(x, x1)
     input_rsg<sup>fb</sup>(y1) \leftarrow sup41(y, y1)
Seed input_rsg<sup>bf</sup>(a) ←
Query query(y) \leftarrow rsg<sup>bf</sup>(a, y)
```



#### Datalog<sup>¬</sup> by example

Accept negative literal in body Complement of transitive closure  $CompG(x,y) \leftarrow \neg G(x,y)$ 

#### More complicated

Some T<sub>P</sub> are not monotone

Some T<sub>P</sub> have no fixpoint containing I

$$- \mathsf{P}_1 = \{\mathsf{p} \leftarrow \neg \mathsf{p}\}$$

$$- \varnothing \to \{\mathsf{p}\} \to \varnothing \to \{\mathsf{p}\} \to \dots$$

- Some T<sub>P</sub> have several minimal fixpoints containing I
  - $P_2 = \{p \leftarrow \neg q, q \leftarrow \neg p\}$
  - Two minimal fixpoints: {p} and {q}.

Some T<sub>P</sub> have a least fixpoint but sequence diverges

- $P_3 = \{p \leftarrow \neg r ; r \leftarrow \neg p; p \leftarrow \neg p, r\}$
- alternates between  $\varnothing$  and {p, r}
- But {p} is a least fixpoint

#### **Model semantics**

- Some programs have no model containing I
- Some program have several minimal models containing

#### First fix: stratification

Impose condition on the syntax

Stratified programs



**Consider more complex semantics** 

- Many such proposals
- Well-founded semantics based on 3-valued logic

# Well-founded by example: 2-player game

e, g are loosing

move graph: (relation K)



There is a pebble in a node

- 2 players alternate playing
- A player moves the pebble following an edge
- A player who cannot move loses

# Winning position

move graph: (relation K)



There is a pebble in a node2 players alternate playingA player moves the pebble following an edgeA player who cannot move loses



There is a pebble in a node2 players alternate playingA player moves the pebble following an edgeA player who cannot move loses

# Program to specify the winning/loosing positions

#### win(x) $\leftarrow$ move(x, y),¬win(y)

Well-founded semantics: find the instance J that agrees with K on move and satisfies the formula corresponding to the rule Instance J – three-valued instance win(d),win(f) are true win(e),win(g) are false win(a),win(b),win(c) are unknown

Fixpoint semantics based on 3-valued logic

4/29/2013

#### **Fixpoint computation**

а

C

yes

win(x) ← move(x, y),¬win(y)
 c b

maybe

- I<sub>0</sub>: win is always false
- I<sub>1</sub>: win: a, b, c, d, f
- I<sub>2</sub>: win: d, f
- I<sub>3</sub>: win: a, b, c, d, f
- I<sub>4</sub>: win: d, f

g

e

No

#### **Complexity and expressivity**

- Datalog and Datalog evaluations are easy
- Datalog⊂ Ptime
  - In the data
  - Inclusion in ptime: polynomial number of stages; each stage in ptime
  - Strict: Expresses only monotone queries;
  - But does not even express all PTIME monotone queries
- Datalog<sup>¬</sup> with well-founded semantics = fixpoint ⊂ Ptime
  - In the data
  - On ordered databases, it is exactly PTIME

# Datalog revival

## Datalog revival

 $\begin{bmatrix} \rho \sigma & \Lambda \end{bmatrix}$ 

Datalog needs to be extended to be useful Updates, value creation, nondeterminism

| Skolem       | [e.g. Gottlob]               |
|--------------|------------------------------|
| Constraints  | [e.g. Revesz]                |
| Time         | [e.g. Chomicki]              |
| Distribution | [e.g. ActiveXML]             |
| Trees        | [e.g. ActiveXML]             |
| Aggregations | [e.g. Consens and Mendelzon] |
| Delegation   | [e.g. Webdamlog]             |

#### Datalog revival: different domains

Declarative networking Data integration and exchange Program verification Data extraction from the Web Knowledge representation

Artifact and workflows Web data management

- [e.g. Lou et al]
- [e.g. Clio, Orchestra]
- [e.g. Semmle]
- [e.g. Gottlob, Lixto]
- [e.g. Gottlob]
- [e.g. ActiveXML] [e.g. Webdamlog]

#### **Declarative networking**

#### Traditional vs. declarative

| Network state    | Distributed database |
|------------------|----------------------|
| Network protocol | Datalog program      |
| Messages         | Messages             |

Series of languages/systems from Hellerstein groups in Berkeley

- Overlog, bloom, dedalus, bud...
- Performance: scalability

Many systems have been developed

Internet routing

**Overlay networks** 

Sensor networks

• • •

#### Data integration

∀Eid, Name, Addr

(employee(Eid, Name, Addr)  $\Rightarrow$ 

∃ Ssn ( name(Ssn, Name) ∧ address(Ssn, Addr) ) )

Use "inverse" rules with Skolem name(ssn(Name, Addr), Name) address(ssn(Name, Addr), Addr)

← employee(X, Name, Addr)

← employee(X, Name, Addr)

Possibly infinite chase and issues with termination

#### Program analysis

Analyze the possible runs of a program

Recursion

Lots of possible runs – lots of data

- Optimization techniques are essential
- Semi-naïve, Magic Sets, Typed-based optimization

#### Data extraction

• Georg's talk next

# Conclusion

50

#### Issues

Give precise semantics to the extensions

Some challenges for the Web

- Scaling to large volumes
- Datalog with distribution
- Datalog with uncertainty
- Datalog with inconsistencies

Berkeley's works

Webdamlog 🖝





