

Advanced Topics in Data Management

Lecture 5

Datalog

Motivation

- SQL can expression *relational queries*;
Iteration/recursion need CTE construct
- Data processing today require iteration.
Common solution: external driver
- Datalog is a language that allows both
recursion and relational queries

Datalog

- Designed in the 80's
- Simple, concise, elegant
- Today is a hot topic: network protocols, static program analysis, DB+ML
- No standard, no reference implementation

Agenda

- Definition
- Naïve Algorithm
- Termination
- Semi-naïve Algorithm

Datalog program

- A datalog program = several rules
- Rules may be recursive
- Set semantics only

Actor(id, fname, lname)
Casts(pid, mid)
Movie(id, name, year)

← Schema

Datalog: Facts and Rules

Actor(id, fname, lname)

Casts(pid, mid)

Movie(id, name, year)

Datalog: Facts and Rules

Facts = tuples in the database

Rules = queries

Actor(id, fname, lname)

Casts(pid, mid)

Movie(id, name, year)

Datalog: Facts and Rules

Facts = tuples in the database

Rules = queries

Actor(344759, 'Douglas', 'Fowley').

Casts(344759, 29851).

Casts(355713, 29000).

Movie(7909, 'A Night in Armour', 1910).

Movie(29000, 'Arizona', 1940).

Movie(29445, 'Ave Maria', 1940).

Actor(id, fname, lname)

Casts(pid, mid)

Movie(id, name, year)

Datalog: Facts and Rules

Facts = tuples in the database

```
Actor(344759, 'Douglas', 'Fowley').  
Casts(344759, 29851).  
Casts(355713, 29000).  
Movie(7909, 'A Night in Armour', 1910).  
Movie(29000, 'Arizona', 1940).  
Movie(29445, 'Ave Maria', 1940).
```

Rules = queries

```
Q1(y) :- Movie(x,y,z), z='1940'.
```

Actor(id, fname, lname)

Casts(pid, mid)

Movie(id, name, year)

Datalog: Facts and Rules

Facts = tuples in the database

```
Actor(344759, 'Douglas', 'Fowley').  
Casts(344759, 29851).  
Casts(355713, 29000).  
Movie(7909, 'A Night in Armour', 1910).  
Movie(29000, 'Arizona', 1940).  
Movie(29445, 'Ave Maria', 1940).
```

Rules = queries

```
Q1(y) :- Movie(x,y,z), z='1940'.
```

Find Movies made in 1940

Actor(id, fname, lname)

Casts(pid, mid)

Movie(id, name, year)

Datalog: Facts and Rules

Facts = tuples in the database

```
Actor(344759, 'Douglas', 'Fowley').  
Casts(344759, 29851).  
Casts(355713, 29000).  
Movie(7909, 'A Night in Armour', 1910).  
Movie(29000, 'Arizona', 1940).  
Movie(29445, 'Ave Maria', 1940).
```

Rules = queries

```
Q1(y) :- Movie(x,y,z), z='1940'.
```

```
Q2(f, l) :- Actor(z,f,l), Casts(z,x),  
           Movie(x,y,'1940').
```

Actor(id, fname, lname)

Casts(pid, mid)

Movie(id, name, year)

Datalog: Facts and Rules

Facts = tuples in the database

Actor(344759, 'Douglas', 'Fowley').
Casts(344759, 29851).
Casts(355713, 29000).
Movie(7909, 'A Night in Armour', 1910).
Movie(29000, 'Arizona', 1940).
Movie(29445, 'Ave Maria', 1940).

Rules = queries

Q1(y) :- Movie(x,y,z), z='1940'.

Q2(f, l) :- Actor(z,f,l), Casts(z,x),
Movie(x,y,'1940').

Find Actors who acted in Movies made in 1940

Actor(id, fname, lname)

Casts(pid, mid)

Movie(id, name, year)

Datalog: Facts and Rules

Facts = tuples in the database

Actor(344759, 'Douglas', 'Fowley').
Casts(344759, 29851).
Casts(355713, 29000).
Movie(7909, 'A Night in Armour', 1910).
Movie(29000, 'Arizona', 1940).
Movie(29445, 'Ave Maria', 1940).

Rules = queries

Q1(y) :- Movie(x,y,z), z='1940'.

Q2(f, l) :- Actor(z,f,l), Casts(z,x),
Movie(x,y,'1940').

Q3(f,l) :- Actor(z,f,l), Casts(z,x1), Movie(x1,y1,1910),
Casts(z,x2), Movie(x2,y2,1940)

Actor(id, fname, lname)

Casts(pid, mid)

Movie(id, name, year)

Datalog: Facts and Rules

Facts = tuples in the database

Actor(344759, 'Douglas', 'Fowley').
Casts(344759, 29851).
Casts(355713, 29000).
Movie(7909, 'A Night in Armour', 1910).
Movie(29000, 'Arizona', 1940).
Movie(29445, 'Ave Maria', 1940).

Rules = queries

Q1(y) :- Movie(x,y,z), z='1940'.

Q2(f, l) :- Actor(z,f,l), Casts(z,x),
Movie(x,y,'1940').

Q3(f,l) :- Actor(z,f,l), Casts(z,x1), Movie(x1,y1,1910),
Casts(z,x2), Movie(x2,y2,1940)

Find Actors who acted in a Movie in 1940 and in one in 1910

Actor(id, fname, lname)

Casts(pid, mid)

Movie(id, name, year)

Datalog: Facts and Rules

Facts = tuples in the database

Actor(344759, 'Douglas', 'Fowley').
Casts(344759, 29851).
Casts(355713, 29000).
Movie(7909, 'A Night in Armour', 1910).
Movie(29000, 'Arizona', 1940).
Movie(29445, 'Ave Maria', 1940).

Rules = queries

Q1(y) :- Movie(x,y,z), z='1940'.

Q2(f, l) :- Actor(z,f,l), Casts(z,x),
Movie(x,y,'1940').

Q3(f,l) :- Actor(z,f,l), Casts(z,x1), Movie(x1,y1,1910),
Casts(z,x2), Movie(x2,y2,1940)

Extensional Database Predicates = EDB = Actor, Casts, Movie

Intensional Database Predicates = IDB = Q1, Q2, Q3

Anatomy of a Rule

Q2(f, l) :- Actor(z,f,l), Casts(z,x), Movie(x,y,'1940').

f, l = head variables

x,y,z = existential variables

Anatomy of a Rule

atom

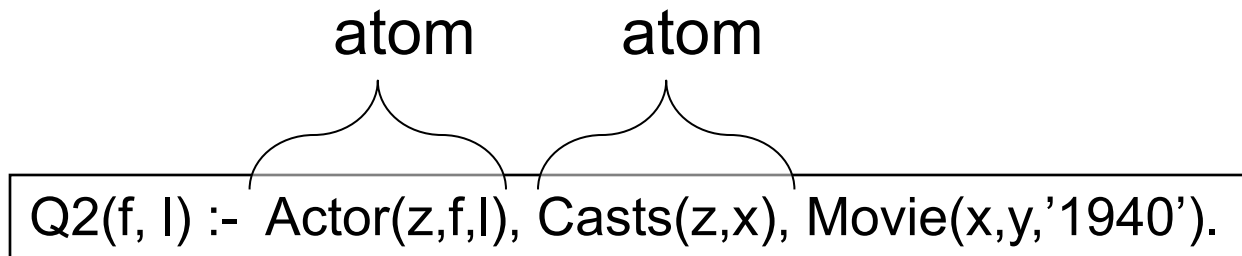


Q2(f, l) :- Actor(z,f,l), Casts(z,x), Movie(x,y,'1940').

f, l = head variables

x,y,z = existential variables

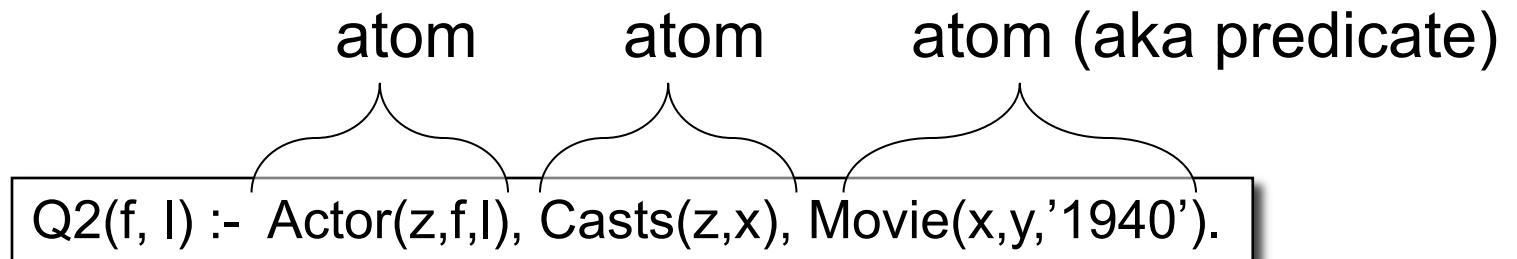
Anatomy of a Rule



f, l = head variables

x, y, z = existential variables

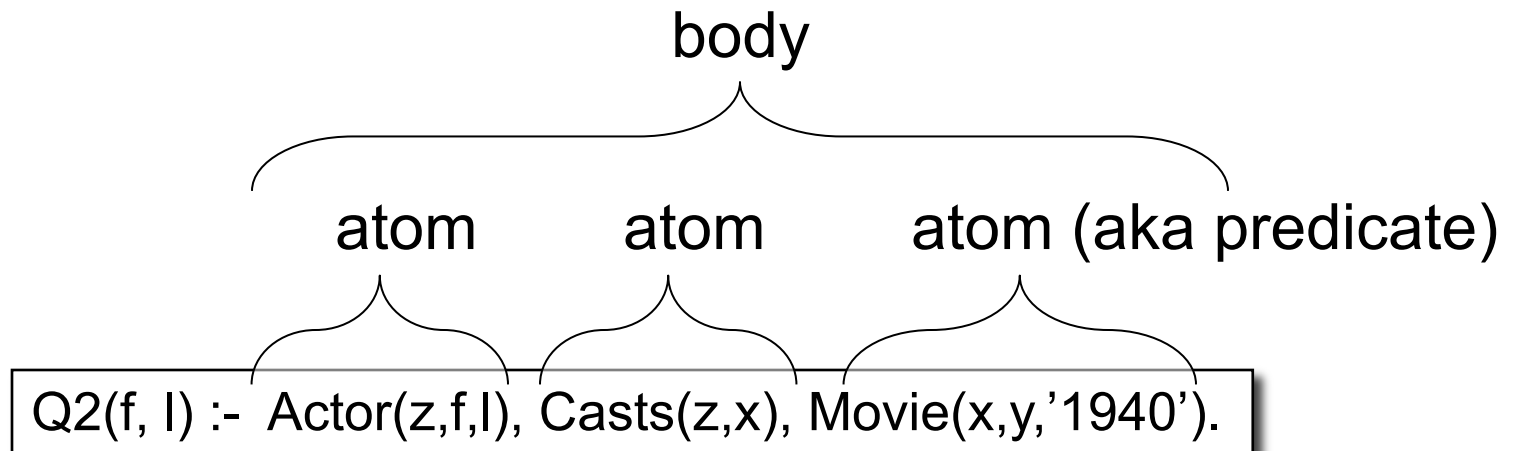
Anatomy of a Rule



f, l = head variables

x,y,z = existential variables

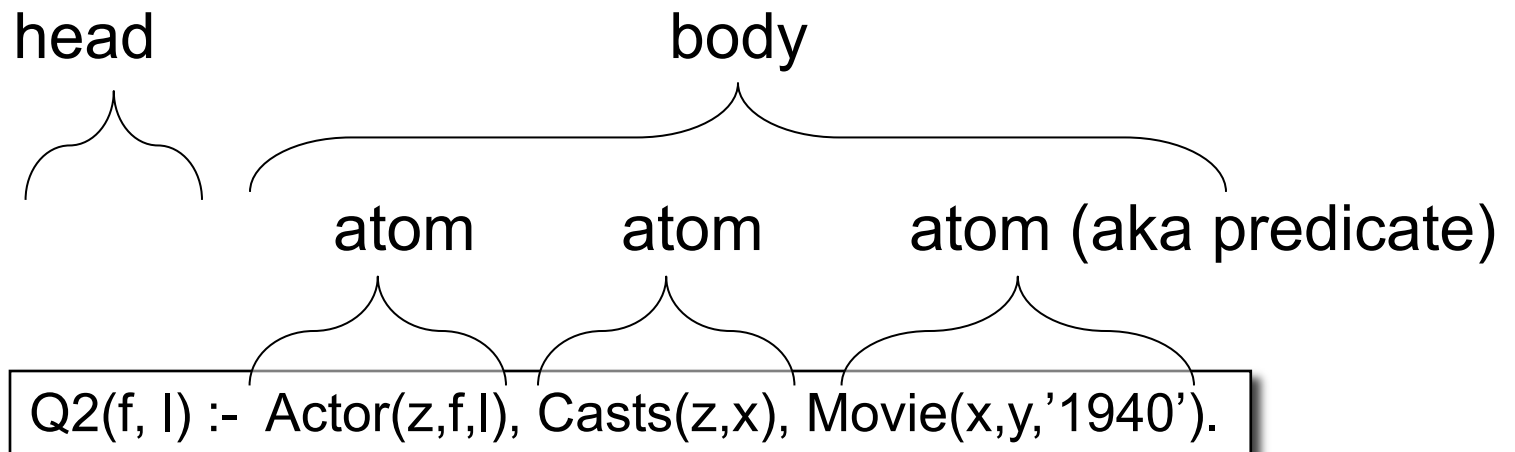
Anatomy of a Rule



f, l = head variables

x, y, z = existential variables

Anatomy of a Rule



f, l = head variables

x, y, z = existential variables

Discussion

- Datalog rules make it very easy to express Selections-Projection-Join (SPJ) queries
 - $Q1(x,y,z) :- R(x,y), S(y,z)$ -- join
 - $Q2(x,y) :- R(x,y), y > 5$ -- selection
 - $Q3(x) :- R(x,y)$ -- projection
 - $Q4(x,z) :- R(x,y), S(y,5), T(y,z)$ -- SPJ

Discussion

- Datalog rules make it very easy to express Selections-Projection-Join (SPJ) queries
 - $Q1(x,y,z) :- R(x,y), S(y,z)$ -- join
 - $Q2(x,y) :- R(x,y), y > 5$ -- selection
 - $Q3(x) :- R(x,y)$ -- projection
 - $Q4(x,z) :- R(x,y), S(y,5), T(y,z)$ -- SPJ
- Unions can be obtained by writing two rules:
 - $Q5(x,y) :- R(x,y)$
 - $Q5(x,y) :- S(x,y)$

Discussion

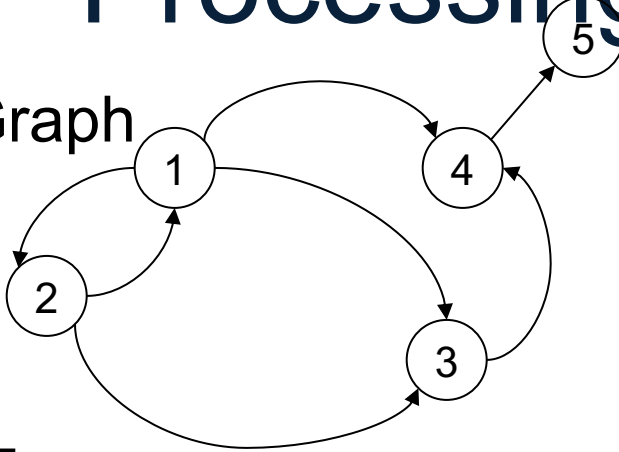
- Datalog rules make it very easy to express Selections-Projection-Join (SPJ) queries
 - $Q1(x,y,z) :- R(x,y), S(y,z)$ -- join
 - $Q2(x,y) :- R(x,y), y > 5$ -- selection
 - $Q3(x) :- R(x,y)$ -- projection
 - $Q4(x,z) :- R(x,y), S(y,5), T(y,z)$ -- SPJ
- Unions can be obtained by writing two rules:
 - $Q5(x,y) :- R(x,y)$
 $Q5(x,y) :- S(x,y)$
- Difference: it's getting complicated; more on this later

Discussion

- Datalog rules make it very easy to express Selections-Projection-Join (SPJ) queries
 - $Q1(x,y,z) :- R(x,y), S(y,z)$ -- join
 - $Q2(x,y) :- R(x,y), y > 5$ -- selection
 - $Q3(x) :- R(x,y)$ -- projection
 - $Q4(x,z) :- R(x,y), S(y,5), T(y,z)$ -- SPJ
- Unions can be obtained by writing two rules:
 - $Q5(x,y) :- R(x,y)$
 $Q5(x,y) :- S(x,y)$
- Difference: it's getting complicated; more on this later
- New! In datalog we can write recursive rules!

Processing Graphs in Datalog

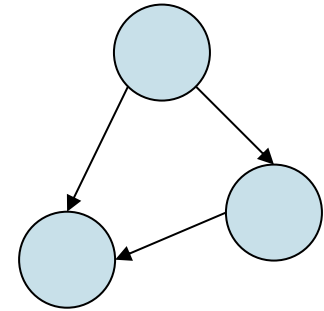
Graph



R=

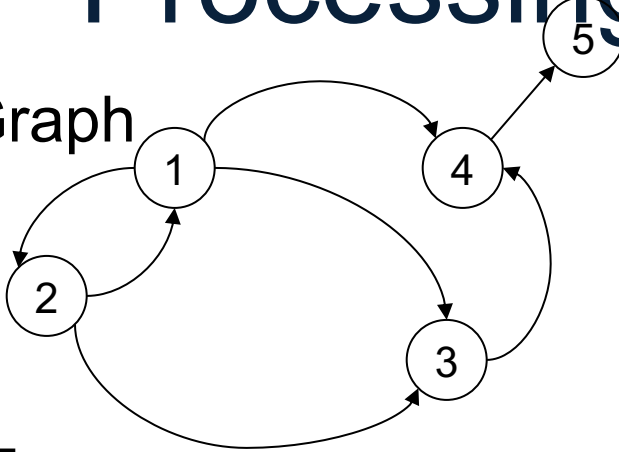
src	dst
1	2
2	1
2	3
1	4
3	4
4	5
1	3

Pattern Matching



Processing Graphs in Datalog

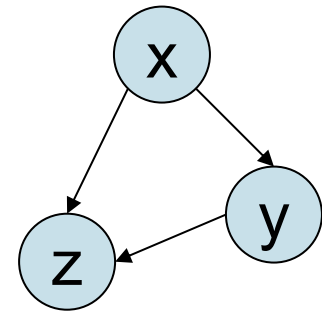
Graph



R=

src	dst
1	2
2	1
2	3
1	4
3	4
4	5
1	3

Pattern Matching



Answer(x,y,z) :- R(x,y), R(x,z), R(y,z)

Discussion

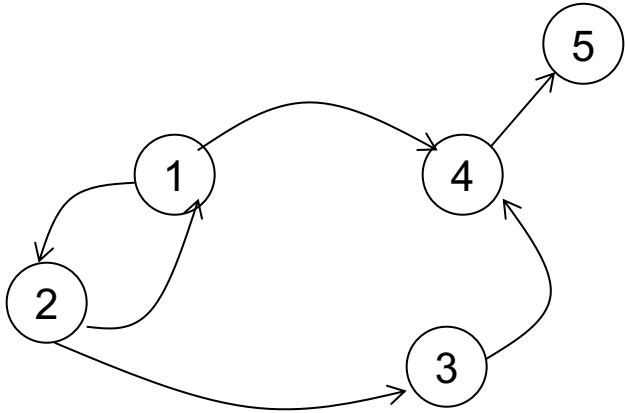
- SQL uses the *named* perspective:
 - Each attribute has a name
 - The order of the attributes is irrelevant
 - **Person.address** or **Person.ssn** or **Person.age**
 - Need to know the names of an attribute

Discussion

- SQL uses the *named* perspective:
 - Each attribute has a name
 - The order of the attributes is irrelevant
 - **Person.address** or **Person.ssn** or **Person.age**
 - Need to know the names of an attribute
- Datalog uses the *unnamed* perspective:
 - The order of the attributes is important
 - Attributes do not have names
 - **Person(x,y,z,u,v)**: here x is ssn, y is address, etc
 - Need to know the position of an attribute

Recursion

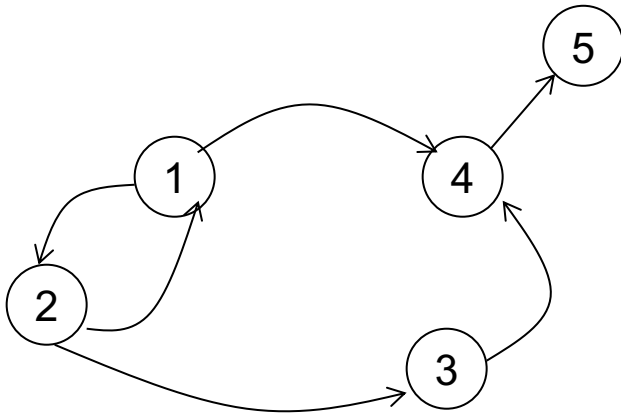
Descendants of node 2



R=

1	2
2	1
2	3
1	4
3	4
4	5

Recursion



R=

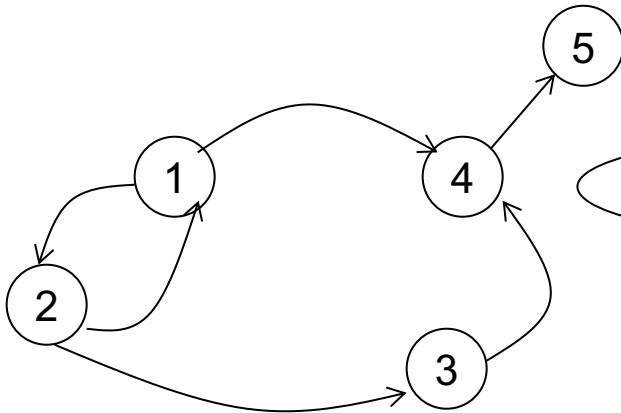
1	2
2	1
2	3
1	4
3	4
4	5

Descendants of node 2

$D(x) :- R(2, x)$

$D(y) :- D(x), R(x, y)$

Recursion



Descendants of node 2

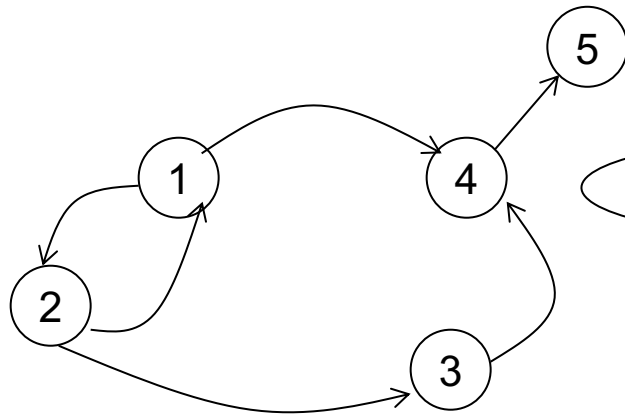
Recursive rule

```
D(x) :- R(2, x)
D(y) :- D(x), R(x,y)
```

R=

1	2
2	1
2	3
1	4
3	4
4	5

Recursion



Descendants of node 2

Recursive rule

```
D(x) :- R(2, x)
D(y) :- D(x), R(x,y)
```

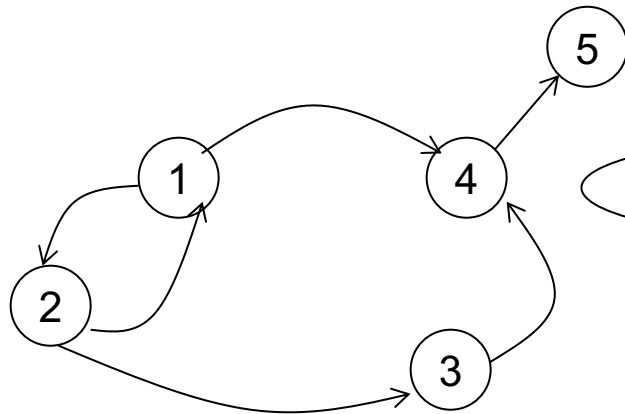
R=

How recursion works in datalog:

Initially D = empty

1	2
2	1
2	3
1	4
3	4
4	5

Recursion



Descendants of node 2

Recursive rule

$D(x) :- R(2, x)$
 $D(y) :- D(x), R(x, y)$

R=

1	2
2	1
2	3
1	4
3	4
4	5

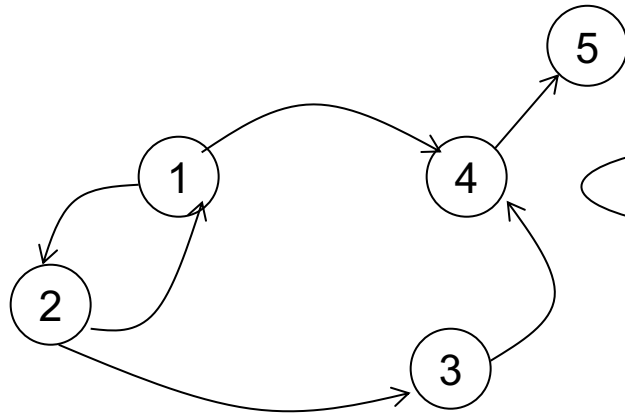
How recursion works in datalog:

Initially D = empty

- Compute both rules:

$D(x) :- R(2, x)$
 $D(y) :- D(x), R(x, y)$

Recursion



R=

1	2
2	1
2	3
1	4
3	4
4	5

How recursion works in datalog:

Initially D = empty

- Compute both rules:
...now D = {1,3}

Descendants of node 2

Recursive rule

$$D(x) \text{ :- } R(2, x)$$

$$D(y) \text{ :- } D(x), R(x, y)$$

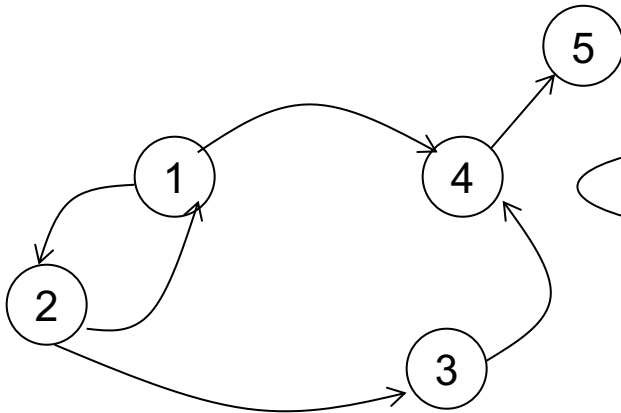
{1,3}

$D(x) \text{ :- } R(2, x)$

{}

$D(y) \text{ :- } D(x), R(x, y)$

Recursion



R=

1	2
2	1
2	3
1	4
3	4
4	5

Recursive rule

Descendants of node 2

$D(x) :- R(2, x)$
 $D(y) :- D(x), R(x, y)$

How recursion works in datalog:

Initially D = empty

- Compute both rules:
...now D = {1,3}
- Compute both rules:

{1,3}

$D(x) :- R(2, x)$

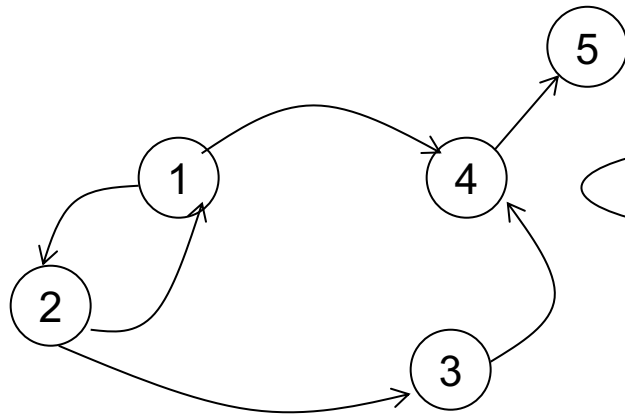
{}

$D(y) :- D(x), R(x, y)$

$D(x) :- R(2, x)$

$D(y) :- D(x), R(x, y)$

Recursion



Recursive rule

Descendants of node 2

$D(x) :- R(2, x)$
 $D(y) :- D(x), R(x, y)$

R=

1	2
2	1
2	3
1	4
3	4
4	5

How recursion works in datalog:

Initially D = empty

- Compute both rules:
...now D = {1,3}
- Compute both rules:
...now D = {1,3,2,4}

{1,3}

$D(x) :- R(2, x)$

{}

$D(y) :- D(x), R(x, y)$

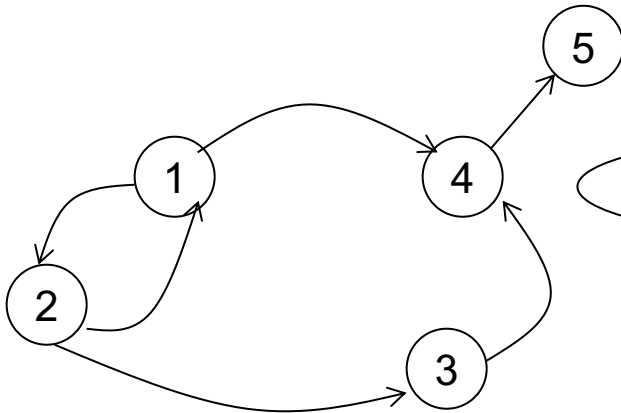
{1,3}

$D(x) :- R(2, x)$

{2,4}

$D(y) :- D(x), R(x, y)$

Recursion



Recursive rule

Descendants of node 2

$D(x) :- R(2, x)$
 $D(y) :- D(x), R(x, y)$

R=

1	2
2	1
2	3
1	4
3	4
4	5

How recursion works in datalog:

Initially D = empty

- Compute both rules:
...now D = {1,3}
- Compute both rules:
...now D = {1,3,2,4}
- Compute both rules:

{1,3}

$D(x) :- R(2, x)$

{}

$D(y) :- D(x), R(x, y)$

{1,3}

$D(x) :- R(2, x)$

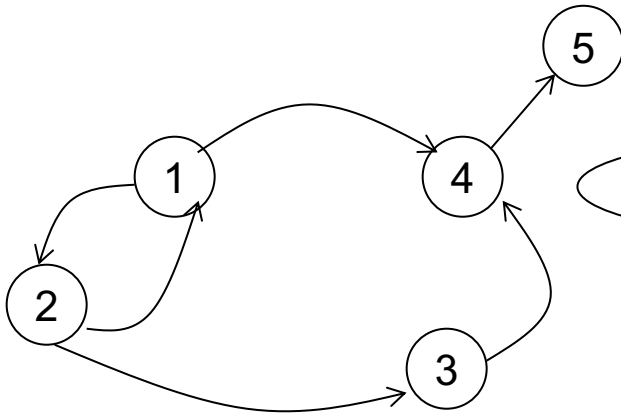
{2,4}

$D(y) :- D(x), R(x, y)$

$D(x) :- R(2, x)$

$D(y) :- D(x), R(x, y)$

Recursion



Recursive rule

Descendants of node 2

$D(x) :- R(2, x)$
 $D(y) :- D(x), R(x,y)$

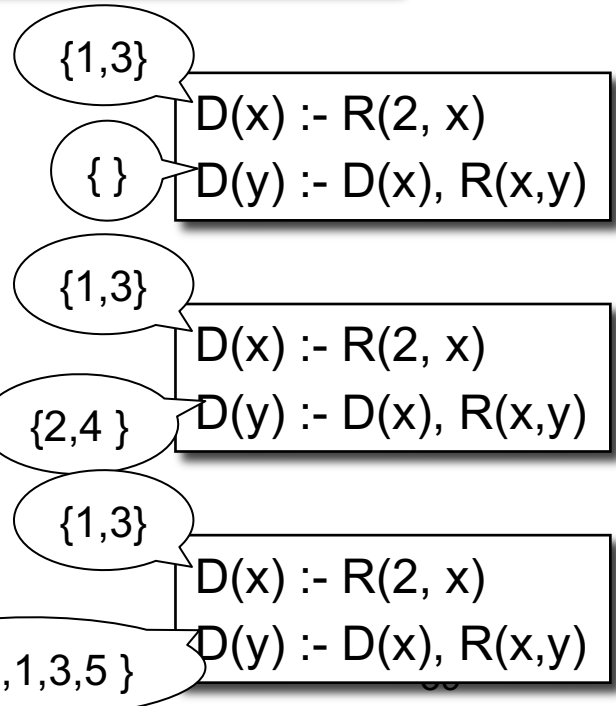
R=

1	2
2	1
2	3
1	4
3	4
4	5

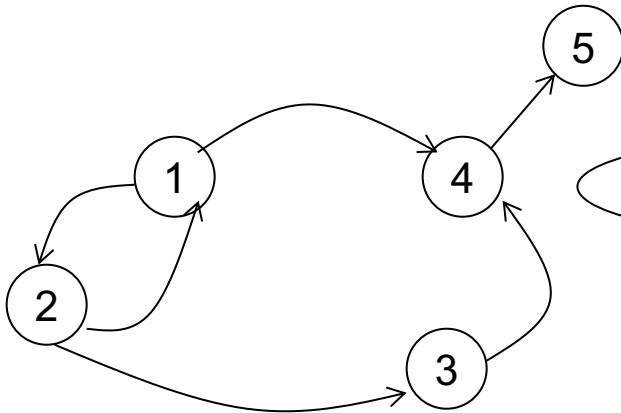
How recursion works in datalog:

Initially D = empty

- Compute both rules:
...now D = {1,3}
- Compute both rules:
...now D = {1,3,2,4}
- Compute both rules:
...now D = {1,3,2,4,5}



Recursion



Recursive rule

Descendants of node 2

$D(x) \text{ :- } R(2, x)$
 $D(y) \text{ :- } D(x), R(x, y)$

R=

1	2
2	1
2	3
1	4
3	4
4	5

How recursion works in datalog:

Initially D = empty

- Compute both rules:
...now D = {1,3}
- Compute both rules:
...now D = {1,3,2,4}
- Compute both rules:
...now D = {1,3,2,4,5}
- Compute both rules:
...nothing new. STOP

{1,3}

$D(x) \text{ :- } R(2, x)$

{}

$D(y) \text{ :- } D(x), R(x, y)$

{1,3}

$D(x) \text{ :- } R(2, x)$

{2,4}

$D(y) \text{ :- } D(x), R(x, y)$

{1,3}

$D(x) \text{ :- } R(2, x)$

{2,4,1,3,5}

$D(y) \text{ :- } D(x), R(x, y)$

Discussion

- Datalog is designed for recursion
- Will prove that it always terminates
- To prove that, first need to define the semantics: **Naïve Algorithm**

Agenda

- Definition
- Naïve Algorithm
- Termination
- Semi-naïve Algorithm

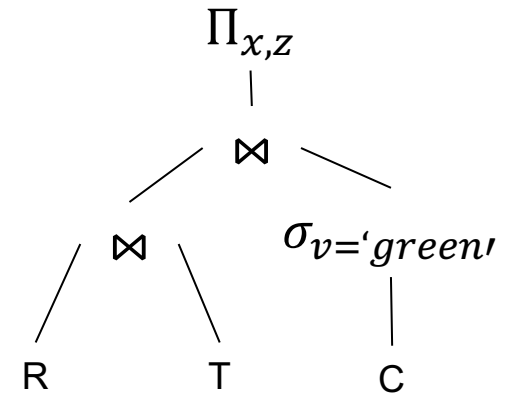
Naïve Evaluation Algorithm

- Convert all datalog rules into a set of Union-Select-Project-Join (USPJ)
- Called: Immediate Consequence Operator
- Naïve Evaluation Algorithm:
Repeatedly apply the **ICO** until fixpoint

Immediate Consequence Operator

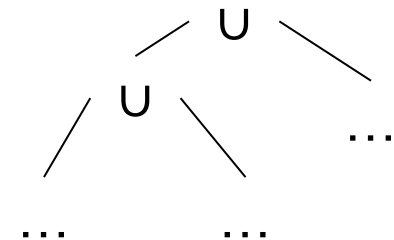
Every rule \rightarrow SPJ query

$T(x,z) :- R(x,y), T(y,z), C(y,'green')$



Multiple rules same IDB head \rightarrow USPJ

$T(x,y) :- \dots$
 $T(x,y) :- \dots$
 \dots



The **ICO** consists of all USPJ's for all IDBs

Naïve Evaluation Algorithm

```
IDB0 := ∅  
for t = 1, ..., ∞ do  
    IDBt := ICO(IDBt-1)  
    if IDBt = IDBt-1 then break
```

Naïve Evaluation Algorithm

$D(x) :- R(2,x)$

$D(y) :- D(x),R(x,y)$

Naïve Evaluation Algorithm

$D(x) :- R(2,x)$

$D(y) :- D(x),R(x,y)$

$\Pi_{R.dst}(\sigma_{R.src=2}(R))$

Naïve Evaluation Algorithm

$D(x) :- R(2,x)$

$D(y) :- D(x),R(x,y)$

$\Pi_{R.dst}(\sigma_{R.src=2}(R)) \cup \Pi_{R.dst}(D \bowtie_{D.node=R.src} R);$

Naïve Evaluation Algorithm

$D(x) :- R(2,x)$

$D(y) :- D(x),R(x,y)$

$\Pi_{R.dst}(\sigma_{R.src=2}(R)) \cup \Pi_{R.dst}(D \bowtie_{D.node=R.src} R);$

Naïve Evaluation Algorithm

$D(x) :- R(2,x)$

$D(y) :- D(x),R(x,y)$

$D := \emptyset;$

repeat

$D := \Pi_{R.dst}(\sigma_{R.src=2}(R)) \cup \Pi_{R.dst}(D \bowtie_{D.node=R.src} R);$

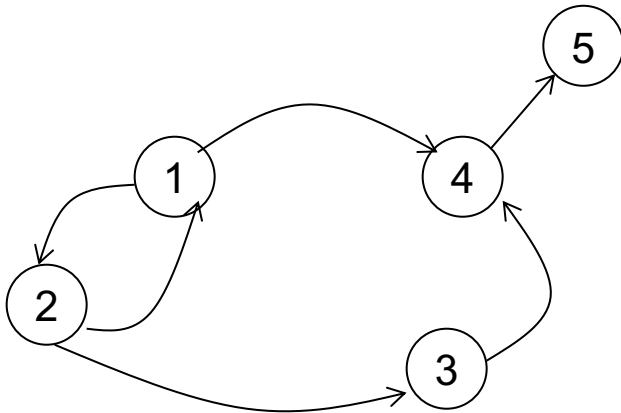
until [no more change]

Naïve Evaluation Algorithm

The Naïve Evaluation Algorithm:

- Always terminates
- Always terminates in a number of steps that is polynomial in the size of the database

Example



R=

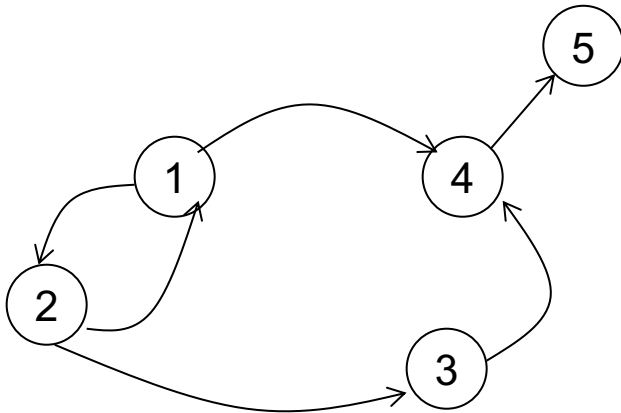
1	2
2	1
2	3
1	4
3	4
4	5

$T(x,y) :- R(x,y)$

$T(x,y) :- R(x,z), T(z,y)$

What does it compute?

Example



R=

1	2
2	1
2	3
1	4
3	4
4	5

Initially:
T is empty.



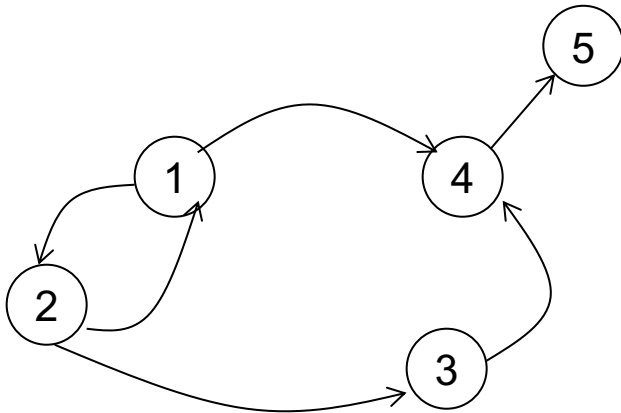
$T(x,y) :- R(x,y)$

$T(x,y) :- R(x,z), T(z,y)$

What does
it compute?

Example

What does it compute?



$T(x,y) :- R(x,y)$
 $T(x,y) :- R(x,z), T(z,y)$

R=

1	2
2	1
2	3
1	4
3	4
4	5

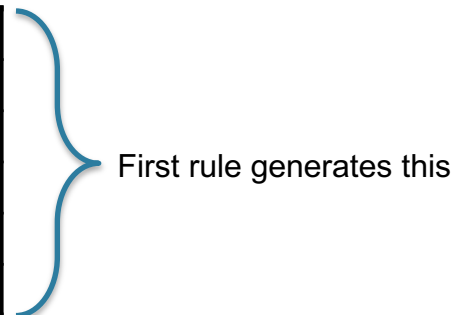
Initially:
T is empty.



First iteration:

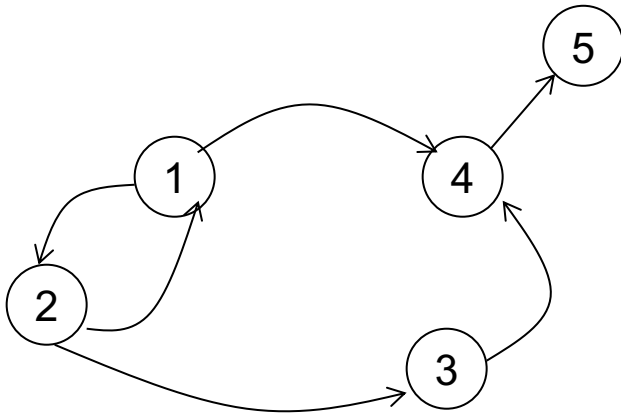
T =

1	2
2	1
2	3
1	4
3	4
4	5



Second rule
generates nothing
(because T is empty)

Example



R=

1	2
2	1
2	3
1	4
3	4
4	5

Initially:
T is empty.



First iteration:
T =

1	2
2	1
2	3
1	4
3	4
4	5

Second iteration:

T =

1	2
2	1
2	3
1	4
3	4
4	5
1	1
2	2
1	3
2	4
1	5
3	5

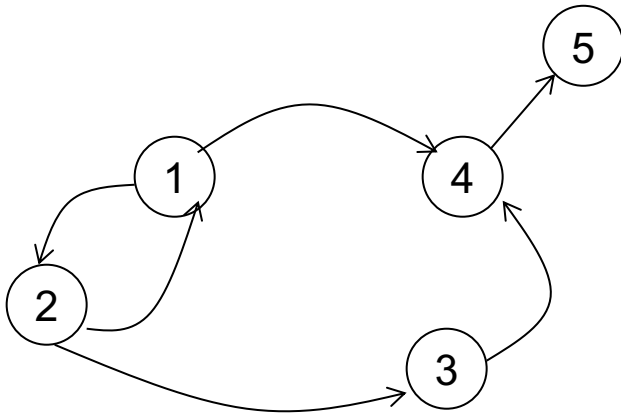
First rule generates this

Second rule generates this

New facts

What does it compute?

Example



R=

1	2
2	1
2	3
1	4
3	4
4	5

Initially:
T is empty.



$T(x,y) \text{ :- } R(x,y)$
 $T(x,y) \text{ :- } R(x,z), T(z,y)$

What does it compute?

First iteration:

T =

1	2
2	1
2	3
1	4
3	4
4	5

Second iteration:

T =

1	2
2	1
2	3
1	4
3	4
4	5
1	1
2	2
1	3
2	4
1	5
3	5

New fact

Third iteration:

T =

1	2
2	1
2	3
1	4
3	4
4	5
1	1
2	2
1	3
2	4
1	5
3	5
2	5

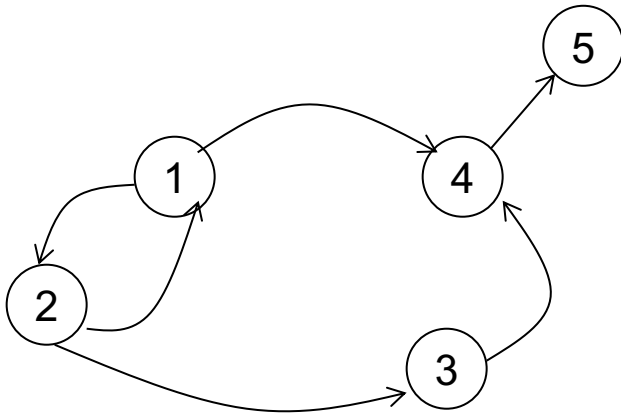
Both rules

First rule

Second rule

56

Example



R =

1	2
2	1
2	3
1	4
3	4
4	5

Initially:
T is empty.

--	--

First iteration:

T =

1	2
2	1
2	3
1	4
3	4
4	5

Second iteration:

T =

1	2
2	1
2	3
1	4
3	4
4	5
1	1
2	2
1	3
2	4
1	5
3	5

Third iteration:

T =

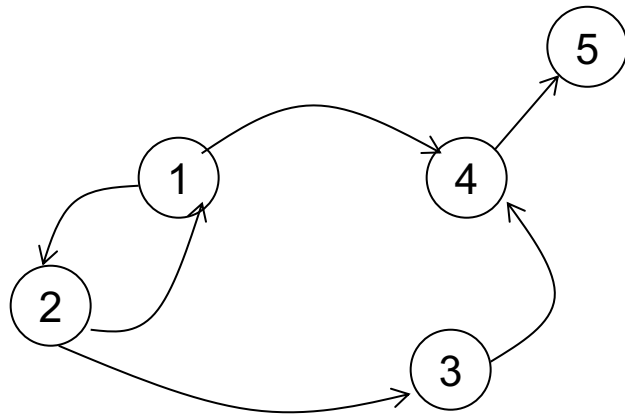
1	2
2	1
2	3
1	4
3	4
4	5
1	1
2	2
1	3
2	4
1	5
3	5
2	5

Fourth iteration
T =
(same)

No new facts.
DONE

What does it compute?

Example



$T(x,y) :- R(x,y)$

$T(x,y) :- R(x,z), T(z,y)$

What does it compute?

R=

1	2
2	1
2	3
1	4
3	4
4	5

Initially:
T is empty.

--	--

First iteration:

T =

1	2
2	1
2	3
1	4
3	4
4	5

Second iteration:

T =

1	2
2	1
2	3
1	4
3	4
4	5
1	1
2	2
1	3
2	4
1	5
3	5

Third iteration:

T =

1	2
2	1
2	3
1	4
3	4
4	5
1	1
2	2
1	3
2	4
1	5
3	5
	5

Fourth iteration
T =
(same)

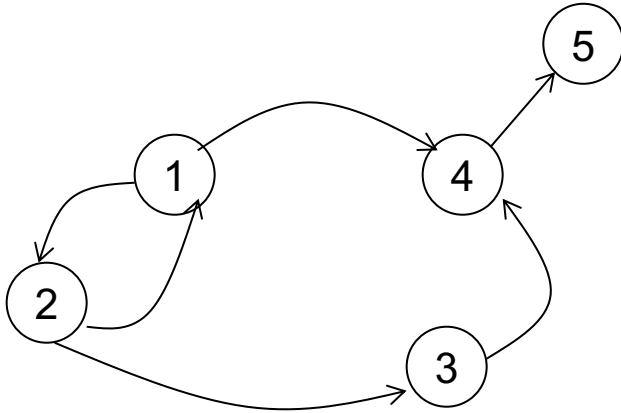
No new facts.
DONE

Datalog

58

Iteration k computes pairs (x,y) connected by path of length $\leq k$

Three Equivalent Programs



R=

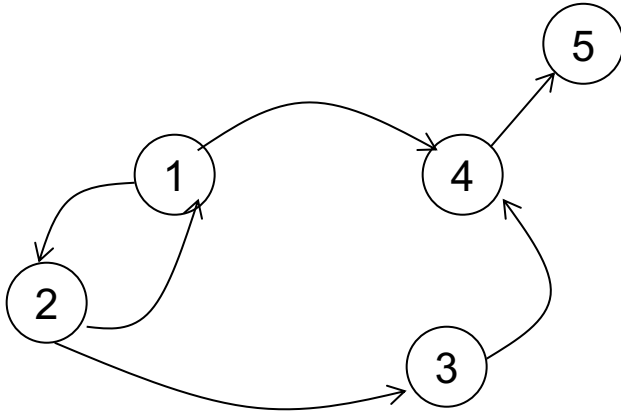
1	2
2	1
2	3
1	4
3	4
4	5

$T(x,y) :- R(x,y)$

$T(x,y) :- R(x,z), T(z,y)$

Right linear

Three Equivalent Programs



R=

1	2
2	1
2	3
1	4
3	4
4	5

$T(x,y) :- R(x,y)$

$T(x,y) :- R(x,z), T(z,y)$

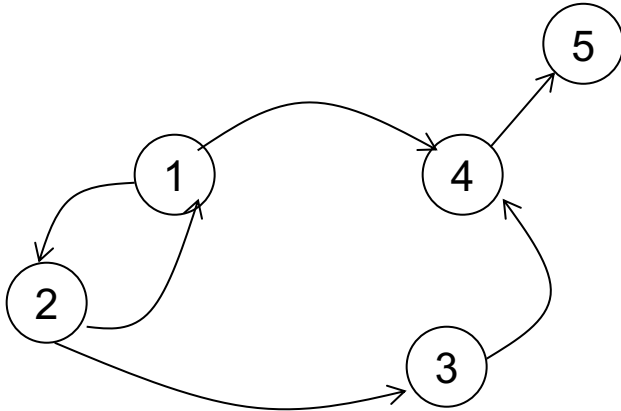
Right linear

$T(x,y) :- R(x,y)$

$T(x,y) :- T(x,z), R(z,y)$

Left linear

Three Equivalent Programs



R=

1	2
2	1
2	3
1	4
3	4
4	5

$T(x,y) :- R(x,y)$

$T(x,y) :- R(x,z), T(z,y)$

Right linear

$T(x,y) :- R(x,y)$

$T(x,y) :- T(x,z), R(z,y)$

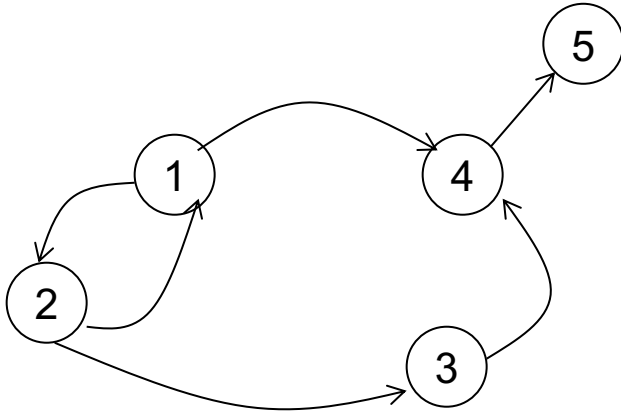
Left linear

$T(x,y) :- R(x,y)$

$T(x,y) :- T(x,z), T(z,y)$

Non-linear

Three Equivalent Programs



R=

1	2
2	1
2	3
1	4
3	4
4	5

$T(x,y) :- R(x,y)$

$T(x,y) :- R(x,z), T(z,y)$

Right linear

$T(x,y) :- R(x,y)$

$T(x,y) :- T(x,z), R(z,y)$

Left linear

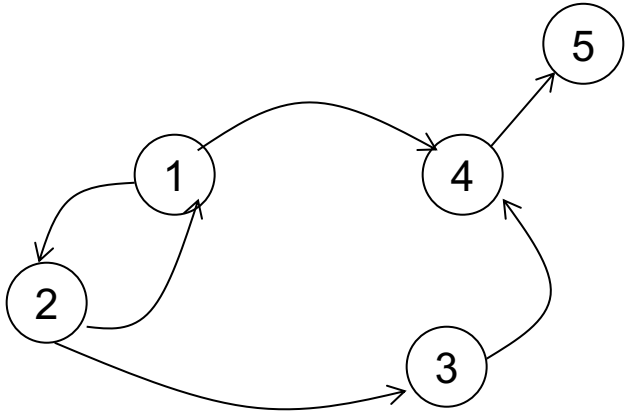
$T(x,y) :- R(x,y)$

$T(x,y) :- T(x,z), T(z,y)$

Non-linear

Question: how many iterations does each require?

Three Equivalent Programs



R=

1	2
2	1
2	3
1	4
3	4
4	5

#iterations = diameter

#iterations = log(diameter)

$T(x,y) :- R(x,y)$
 $T(x,y) :- R(x,z), T(z,y)$

Right linear

$T(x,y) :- R(x,y)$
 $T(x,y) :- T(x,z), R(z,y)$

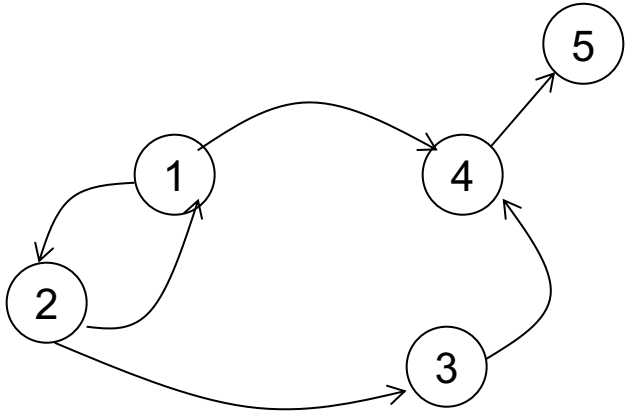
Left linear

$T(x,y) :- R(x,y)$
 $T(x,y) :- T(x,z), T(z,y)$

Non-linear

Question: how many iterations does each require?

Multiple IDBs



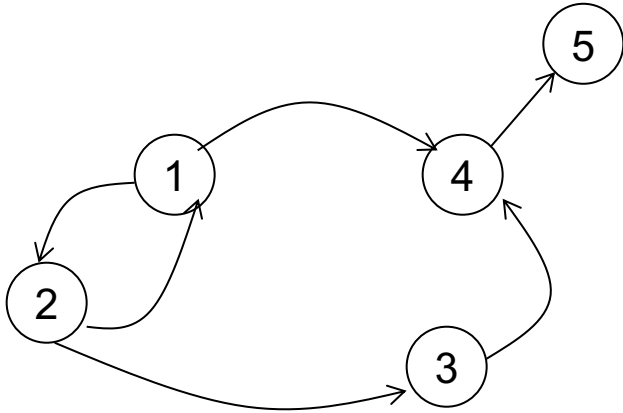
Find pairs of nodes (x,y) connected by a path of even length

R=

1	2
2	1
2	3
1	4
3	4
4	5



Multiple IDBs



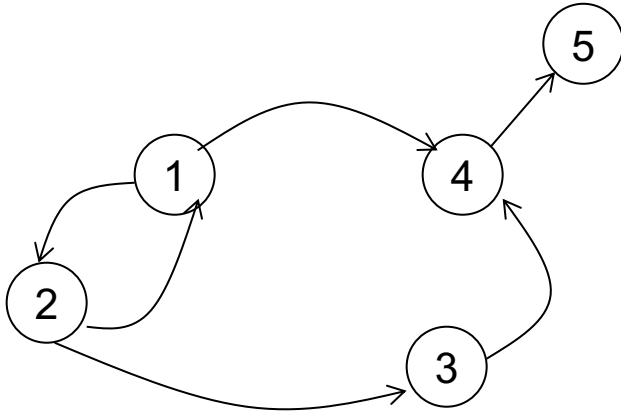
Find pairs of nodes (x,y) connected by a path of even length

R=

1	2
2	1
2	3
1	4
3	4
4	5

Odd(x,y) :- R(x,y)

Multiple IDBs



Find pairs of nodes (x,y)
connected by a path of even length

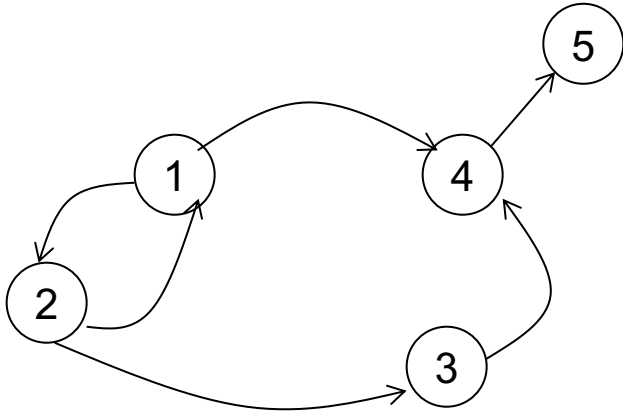
R=

1	2
2	1
2	3
1	4
3	4
4	5

Odd $(x,y) :- R(x,y)$

Even $(x,y) :- Odd(x,z), R(z,y)$

Multiple IDBs



Find pairs of nodes (x,y)
connected by a path of even length

R=

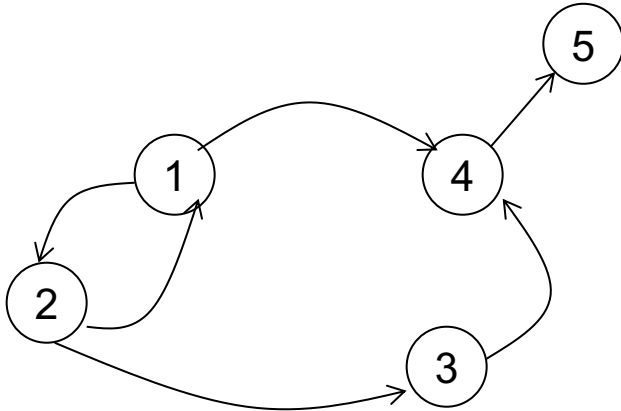
1	2
2	1
2	3
1	4
3	4
4	5

Odd $(x,y) :- R(x,y)$

Even $(x,y) :- Odd(x,z), R(z,y)$

Odd $(x,y) :- Even(x,z), R(z,y)$

Multiple IDBs



Find pairs of nodes (x,y) connected by a path of even length

R=

1	2
2	1
2	3
1	4
3	4
4	5

Odd $(x,y) :- R(x,y)$

Even $(x,y) :- Odd(x,z), R(z,y)$

Odd $(x,y) :- Even(x,z), R(z,y)$

We have two IDBs: Odd (x,y) and Even (x,y)

Naïve Evaluation Algorithm

When multiple IDBs:
need to compute their
new values together:

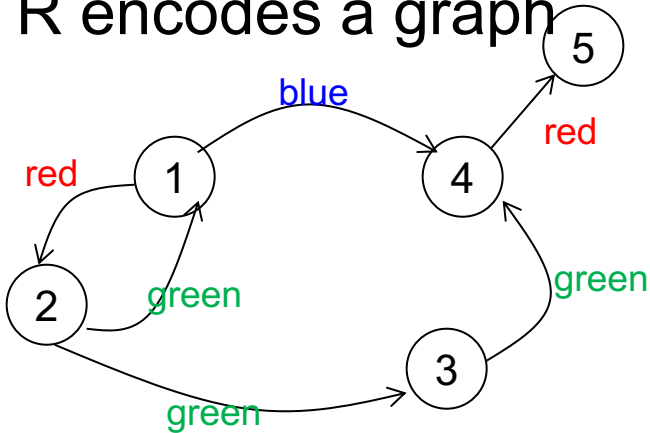
```
Odd(x,y)  :- R(x,y)
Even(x,y) :- Odd(x,z),R(z,y)
Odd(x,y)  :- Even(x,z),R(z,y)
```

ICO

```
Odd :=  $\emptyset$ ; Even :=  $\emptyset$ ;
repeat
  Evennew :=  $\Pi_{x,y}(\text{Odd} \bowtie R)$ ;
  Oddnew :=  $R \cup \Pi_{x,y}(\text{Even} \bowtie R)$ ;
  if Odd=Oddnew  $\wedge$  Even=Evennew
    then break
  Odd:=Oddnew
  Even:=Evennew
```

Labeled Graphs

R encodes a graph



R=

1	2	red
2	1	green
2	3	green
1	4	blue
3	4	green
4	5	red

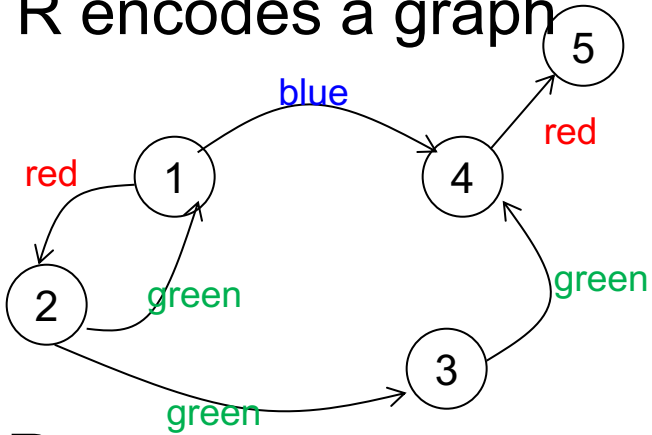
Find pairs of nodes (x,y) connected by a green path

GreenP(x,y) :- R(x,y,'green')

GreenP(x,y) :- R(x,z,'green'),GreenP(z,y)

Labeled Graphs

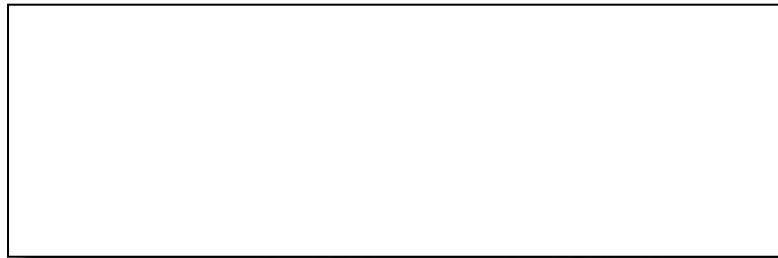
R encodes a graph



R=

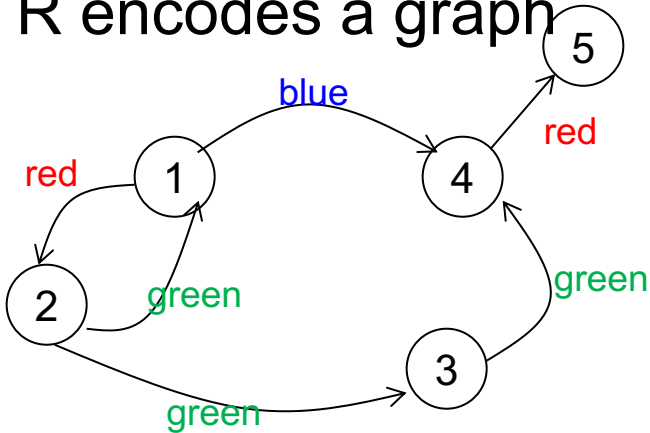
1	2	red
2	1	green
2	3	green
1	4	blue
3	4	green
4	5	red

Find pairs of nodes (x,y) connected by a monochromatic path



Labeled Graphs

R encodes a graph



R=

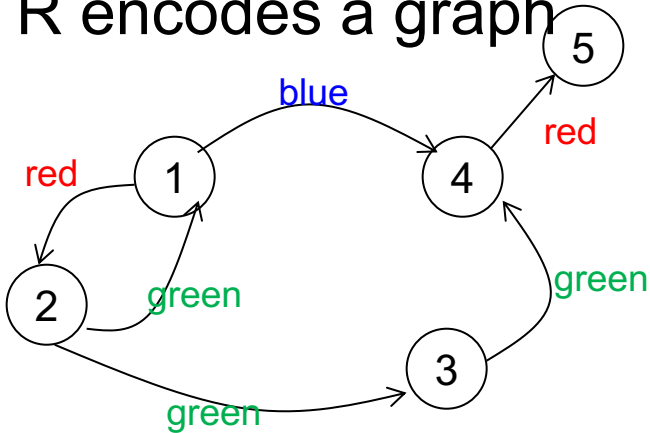
1	2	red
2	1	green
2	3	green
1	4	blue
3	4	green
4	5	red

Find pairs of nodes (x,y) connected by a monochromatic path

$P(x,y,c) :- R(x,y,c)$

Labeled Graphs

R encodes a graph



R=

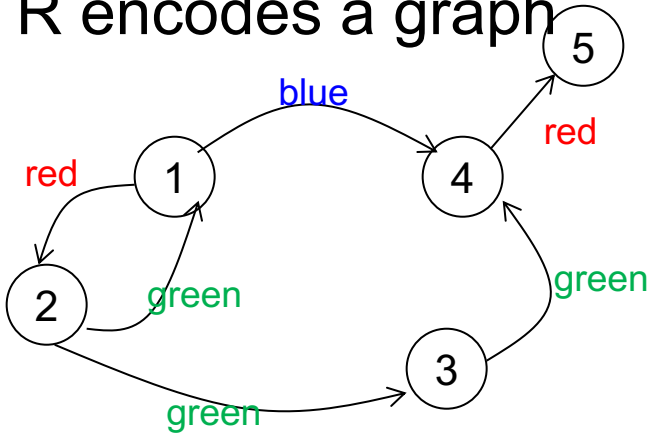
1	2	red
2	1	green
2	3	green
1	4	blue
3	4	green
4	5	red

Find pairs of nodes (x,y) connected by a monochromatic path

$$P(x,y,c) \text{ :- } R(x,y,c)$$
$$P(x,y,c) \text{ :- } R(x,z,c), P(z,y,c)$$

Labeled Graphs

R encodes a graph



R=

1	2	red
2	1	green
2	3	green
1	4	blue
3	4	green
4	5	red

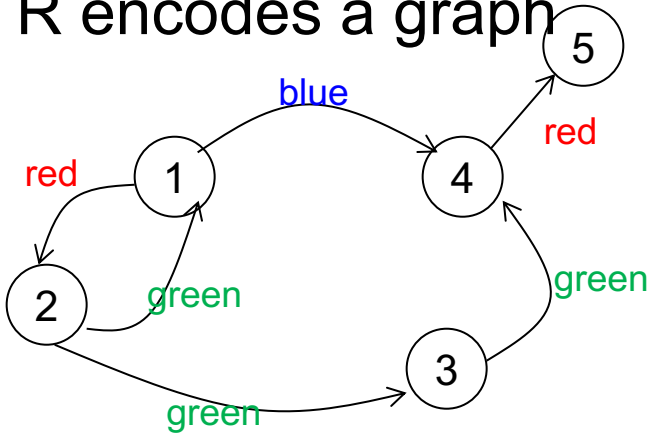
Find pairs of nodes (x,y)
connected by a monochromatic path

We join on
both the node z,
and the color c

$$P(x,y,c) \text{ :- } R(x,y,c)$$
$$P(x,y,c) \text{ :- } R(x,z,c), P(z,y,c)$$

Labeled Graphs

R encodes a graph



R=

1	2	red
2	1	green
2	3	green
1	4	blue
3	4	green
4	5	red

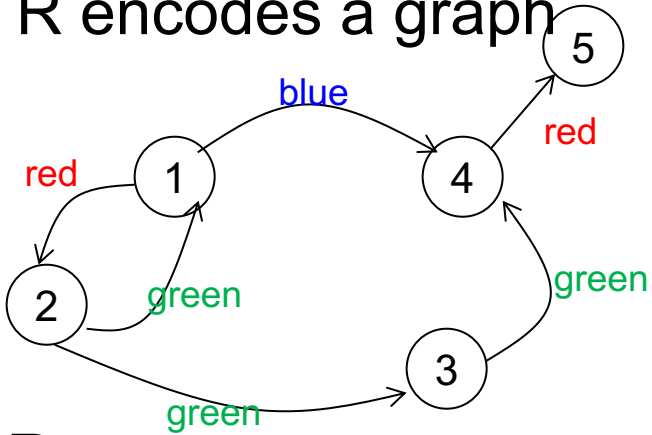
Find pairs of nodes (x,y)
connected by a monochromatic path

We join on
both the node z,
and the color c

```
P(x,y,c) :- R(x,y,c)
P(x,y,c) :- R(x,z,c),P(z,y,c)
Answer(x,y) :- P(x,y,c)
```

Labeled Graphs

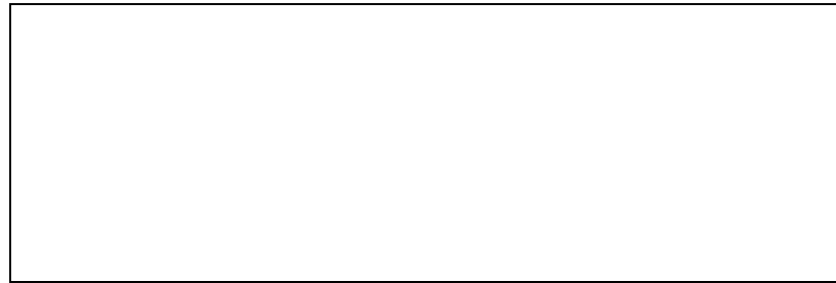
R encodes a graph



R=

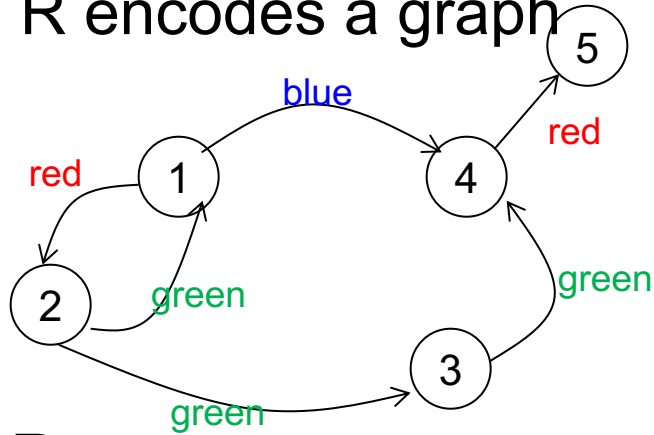
1	2	red
2	1	green
2	3	green
1	4	blue
3	4	green
4	5	red

Find all nodes reachable from node 2 by a path containing exactly one red edge.



Labeled Graphs

R encodes a graph

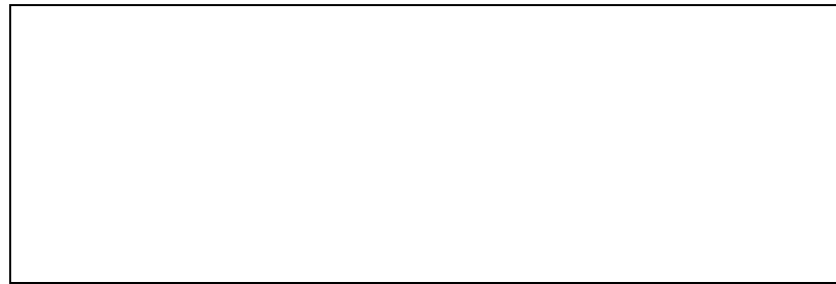
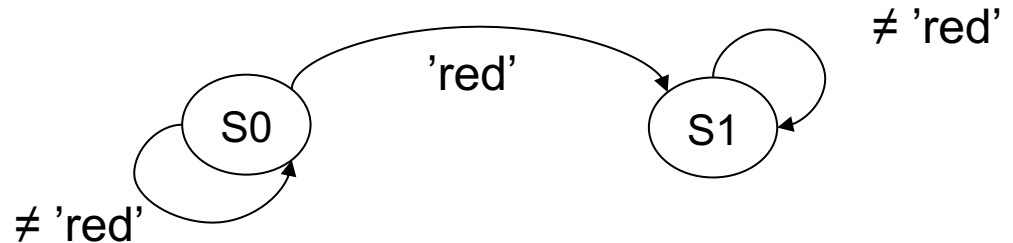


R=

1	2	red
2	1	green
2	3	green
1	4	blue
3	4	green
4	5	red

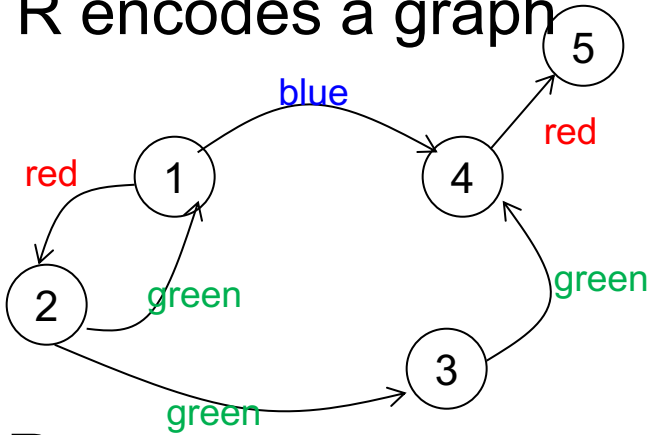
Find all nodes reachable from node 2 by a path containing exactly one red edge.

Automaton:



Labeled Graphs

R encodes a graph

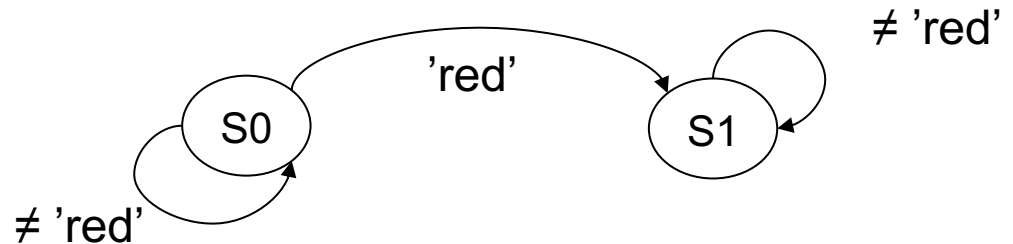


R=

1	2	red
2	1	green
2	3	green
1	4	blue
3	4	green
4	5	red

Find all nodes reachable from node 2 by a path containing exactly one red edge.

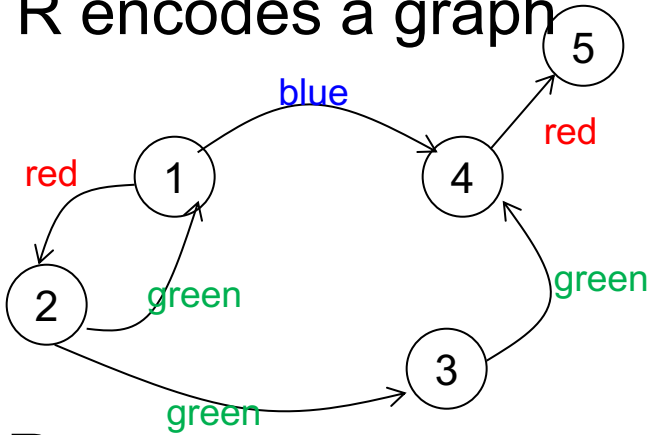
Automaton:



S0(2) :- .

Labeled Graphs

R encodes a graph

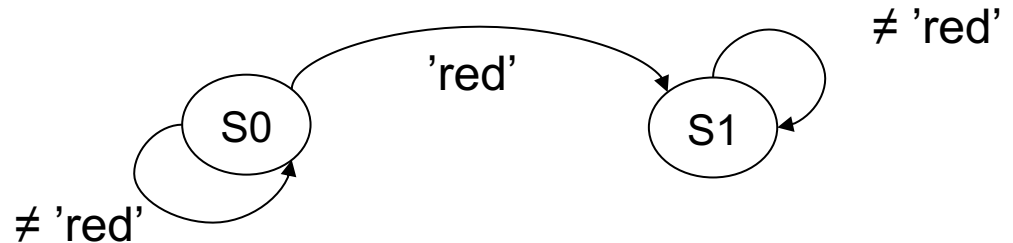


R=

1	2	red
2	1	green
2	3	green
1	4	blue
3	4	green
4	5	red

Find all nodes reachable from node 2 by a path containing exactly one red edge.

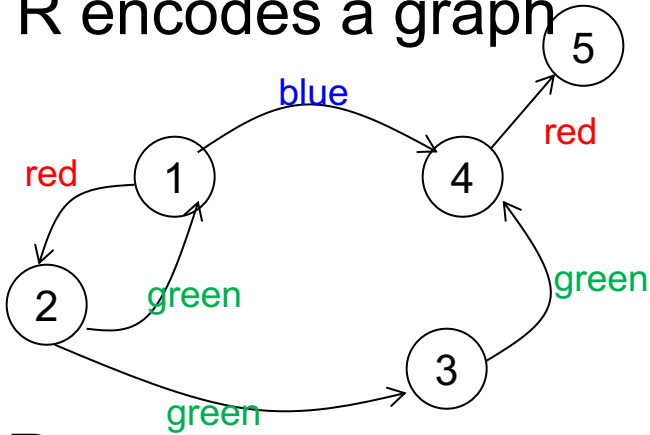
Automaton:



$S0(2)$	$:- .$
$S0(y)$	$:- S0(x), R(x,y,c), c \neq \text{'red'}$.

Labeled Graphs

R encodes a graph

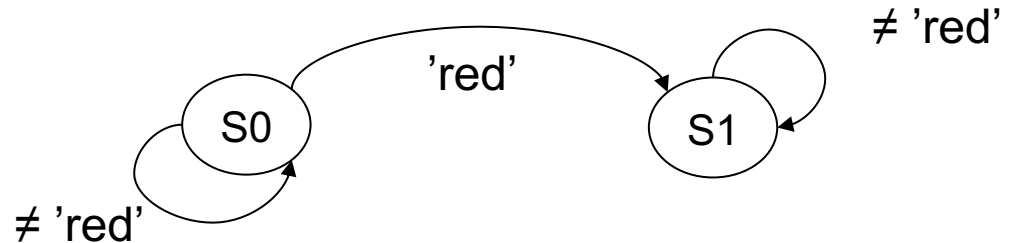


R=

1	2	red
2	1	green
2	3	green
1	4	blue
3	4	green
4	5	red

Find all nodes reachable from node 2 by a path containing exactly one red edge.

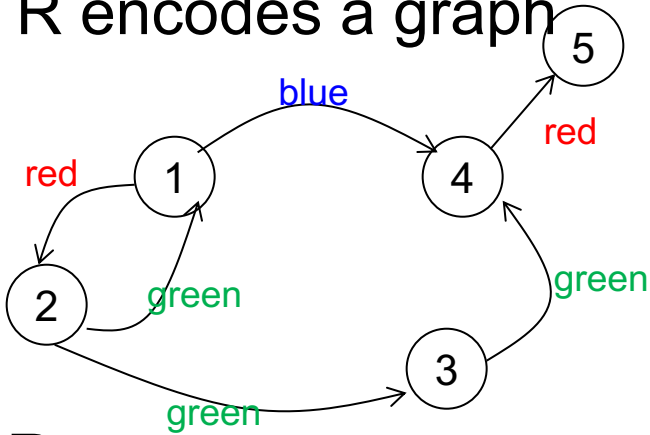
Automaton:



S0(2)	:- .
S0(y)	:- S0(x), R(x,y,c), c!='red'.
S1(y)	:- S0(x), R(x,y,'red').

Labeled Graphs

R encodes a graph

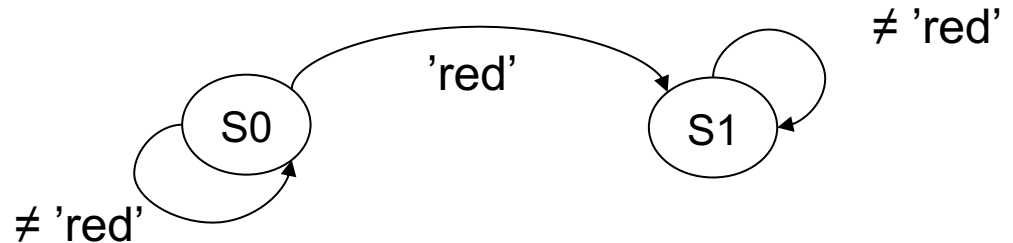


R=

1	2	red
2	1	green
2	3	green
1	4	blue
3	4	green
4	5	red

Find all nodes reachable from node 2 by a path containing exactly one red edge.

Automaton:



S0(2)	:- .
S0(y)	:- S0(x), R(x,y,c), c!='red'.
S1(y)	:- S0(x), R(x,y,'red').
S1(y)	:- S1(x), R(x,y,c), c!='red'.

Discussion: Recursion in SQL

SQL has limited form of recursion, BUT:

Discussion: Recursion in SQL

SQL has limited form of recursion, BUT:

- Single IDB
 - Called: Common Table Expression, CTE
 - Cannot write Odd/Even, Red/NoRed, etc

Discussion: Recursion in SQL

SQL has limited form of recursion, BUT:

- Single IDB
 - Called: Common Table Expression, CTE
 - Cannot write Odd/Even, Red/NoRed, etc
- Linear query only
 - Cannot write $T(x,y) :- T(x,z), T(z,y)$

Discussion: Recursion in SQL

SQL has limited form of recursion, BUT:

- Single IDB
 - Called: Common Table Expression, CTE
 - Cannot write Odd/Even, Red/NoRed, etc
- Linear query only
 - Cannot write $T(x,y) :- T(x,z), T(z,y)$
- Has bag semantics (really???)
 - May not terminate!

Discussion: Recursion in SQL

```
T(x,y) :- R(x,y)
```

```
T(x,y) :- R(x,z), T(z,y)
```

Discussion: Recursion in SQL

```
T(x,y) :- R(x,y)
```

```
T(x,y) :- R(x,z), T(z,y)
```

```
with recursive T as(  
  select * from R  
  union  
  select distinct R.x, T.y  
  from R, T  
  where R.y=T.x  
)  
select * from T;
```

Discussion: Recursion in SQL

```
T(x,y) :- R(x,y)
T(x,y) :- R(x,z), T(z,y)
```

Relation T is called a
Common Table Expression CTE

```
with recursive T as(
  select * from R
  union
  select distinct R.x, T.y
  from R, T
  where R.y=T.x
)
select * from T;
```


Agenda

- Definition
- Naïve Algorithm
- Termination
- Semi-naïve Algorithm

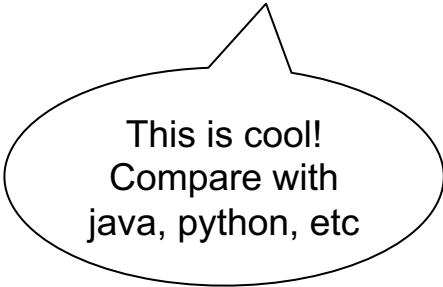
Runtime of a Datalog Program

Theorem The Naïve Algorithm:

- Always terminates
- Terminates in a number of steps that is polynomial in the size of the database

Assumptions:

- Set semantics only
- Monotone rules only
- No “value invention”

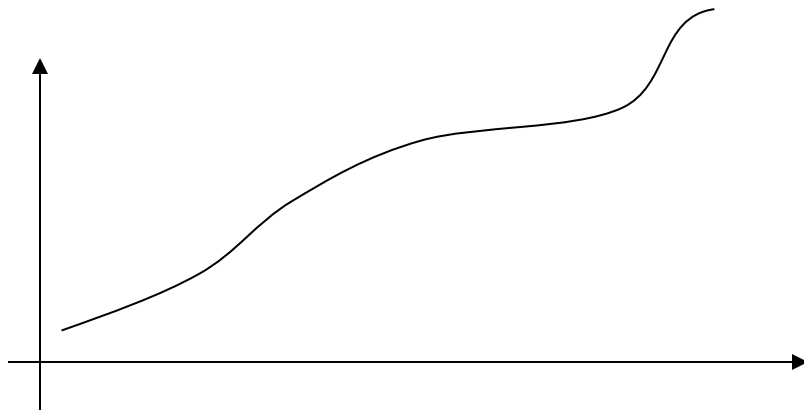


This is cool!
Compare with
java, python, etc

Detour: Monotone Functions

- A function $f(x)$ is called monotonically increasing, or just monotone if:

If $x \leq y$ then $f(x) \leq f(y)$



Monotone Queries

- A query with input relations R, S, T, \dots is called monotone if, whenever we increase a relation, the query answer also increases (or stays the same)
- Increase here means larger set

Monotone Queries

- A query with input relations R, S, T, \dots is called monotone if, whenever we increase a relation, the query answer also increases (or stays the same)
- Increase here means larger set
- Mathematically

If $R \subseteq R', S \subseteq S', \dots$ then $Q(R, S, \dots) \subseteq Q(R', S', \dots)$

Which Ops are Monotone?

- Selection: σ_{pred}
- Projection: $\Pi_{A,B,\dots}$
- Join: \bowtie
- Union: \cup
- Difference: $-$
- Group-by-sum: $\gamma_{A,B,sum}(C)$

Which Ops are Monotone?

- Selection: σ_{pred} **MONOTONE**
- Projection: $\Pi_{A,B,\dots}$ **MONOTONE**
- Join: \bowtie **MONOTONE**
- Union: \cup **MONOTONE**
- Difference: $-$ **NON-MONOTONE**
- Group-by-sum: $\gamma_{A,B,sum}(C)$ **NON-MONOTONE**

Back to Datalog

Naïve Algorithm:

- Always terminates
- Terminates in a number of steps that is polynomial in the size of the database
- This is cool!
Compare with java, python, etc

Assumptions:

- Set semantics only
- Monotone rules only
- No “value invention”

Will show this next


```
IDB0 := ∅; t := 0
```

```
repeat IDBt+1 := USPJ(IDBt); t := t + 1
```

```
until no more change
```

Proof

Fact: every USPJ query is monotone

Proof: uses only σ , Π , \bowtie , \cup

$IDB_0 := \emptyset; \quad t := 0$

repeat $IDB_{t+1} := USPJ(IDB_t); \quad t := t + 1$

until no more change

Proof

Fact: every USPJ query is monotone

Proof: uses only σ, Π, \bowtie, U

Fact: the IDBs increase: $IDB_t \subseteq IDB_{t+1}$

Proof: by induction

$IDB_0 := \emptyset; \quad t := 0$

repeat $IDB_{t+1} := USPJ(IDB_t); \quad t := t + 1$

until no more change

Proof

Fact: every USPJ query is monotone

Proof: uses only $\sigma, \Pi, \bowtie, \cup$

Fact: the IDBs increase: $IDB_t \subseteq IDB_{t+1}$

Proof: by induction $IDB_0 (= \emptyset) \subseteq IDB_1$

```
IDB0 := ∅; t := 0
repeat IDBt+1 := USPJ(IDBt); t := t + 1
until no more change
```

Proof

Fact: every USPJ query is monotone

Proof: uses only σ , Π , \bowtie , \cup

Fact: the IDBs increase: $IDB_t \subseteq IDB_{t+1}$

Proof: by induction $IDB_0 (= \emptyset) \subseteq IDB_1$

Assuming $IDB_t \subseteq IDB_{t+1}$ we have:
 $USPJ(IDB_t) \subseteq USPJ(IDB_{t+1})$

$IDB_0 := \emptyset; \quad t := 0$

repeat $IDB_{t+1} := USPJ(IDB_t); \quad t := t + 1$

until no more change

Proof

Fact: every USPJ query is monotone

Proof: uses only $\sigma, \Pi, \bowtie, \cup$

Fact: the IDBs increase: $IDB_t \subseteq IDB_{t+1}$

Proof: by induction $IDB_0 (= \emptyset) \subseteq IDB_1$

Assuming $IDB_t \subseteq IDB_{t+1}$ we have:

$$IDB_{t+1} = USPJ(IDB_t) \subseteq USPJ(IDB_{t+1}) = IDB_{t+2}$$

Naïve Evaluation Algorithm

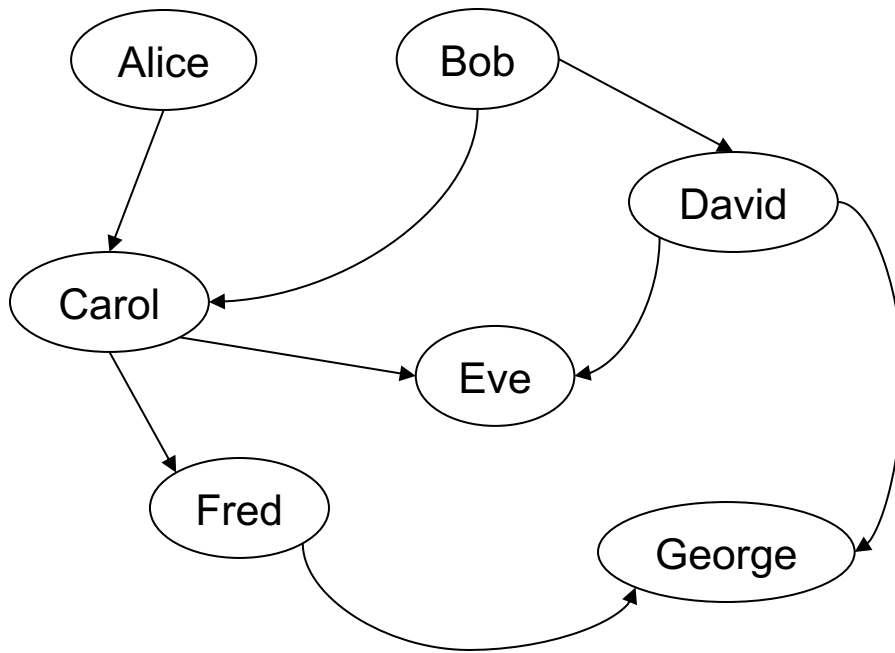
Consequence: The naïve algorithm terminates, in $O(n^k)$ steps, where:

- n = number of distinct values in the DB
- k = arity of widest IDB relation

Proof: IDBs increases to $\leq O(n^k)$ facts

Same Generation

Two people are in the same generation if they are descendants at the same generation of some common ancestor



SG

p1	p2
Carol	David
Eve	George
Fred	George
Fred	Eve

Same Generation

Compute pairs of people at the same generation

```
// common parent
```


Same Generation

Compute pairs of people at the same generation

```
// common parent  
SG(x,y) :- ParentChild(p,x), ParentChild(p,y)
```

Same Generation

Compute pairs of people at the same generation

```
// common parent  
SG(x,y) :- ParentChild(p,x), ParentChild(p,y)  
  
// parents at the same generation
```

Same Generation

Compute pairs of people at the same generation

```
// common parent
```

```
SG(x,y) :- ParentChild(p,x), ParentChild(p,y)
```

```
// parents at the same generation
```

```
SG(x,y) :- ParentChild(p,x), ParentChild(q,y), SG(p,q)
```

Same Generation

Compute pairs of people at the same generation

```
// common parent
SG(x,y) :- ParentChild(p,x), ParentChild(p,y)

// parents at the same generation
SG(x,y) :- ParentChild(p,x), ParentChild(q,y), SG(p,q)
```

Problem: this includes answers like SG(Carol, Carol)

And also SG(Eve, George), SG(George, Eve)

How to fix?

Same Generation

Compute pairs of people at the same generation

```
// common parent
```

```
SG(x,y) :- ParentChild(p,x), ParentChild(p,y),  $x < y$ 
```

```
// parents at the same generation
```

```
SG(x,y) :- ParentChild(p,x), ParentChild(q,y),  
           SG(p,q),  $x < y$ 
```

Agenda

- Definition
- Naïve Algorithm
- Termination
- Semi-naïve Algorithm

Main Idea

- The naïve algorithm re-re-computes the same facts over and over again
- Main idea: compute only the new facts of the IDBs, from the previous new facts
- Need: Incremental View Maintenance

Incremental View Maintenance

- A view is a relation defined by a query

```
CREATE VIEW PathsLengthTwo AS
  SELECT e1.src as src, e2.dst as dst
  FROM Edge e1, Edge e2
  WHERE e1.dst = e2.src
```

- In datalog: every IDB is a view

```
PathsLengthTwo(x,y) :- Edge(x,z),Edge(z,y)
```


Incremental View Maintenance

- Suppose we have computed the view:
 $V := [\text{monotone query here}]$

Incremental View Maintenance

- Suppose we have computed the view:
 $V := [\text{monotone query here}]$
- One, or several of the base relations is updated: e.g. we insert new tuples:

$$R := R \cup \Delta R$$

Incremental View Maintenance

- Suppose we have computed the view:

$$V := [\text{monotone query here}]$$

- One, or several of the base relations is updated: e.g. we insert new tuples:

$$R := R \cup \Delta R$$

- Then the view needs to be incremented with some new tuples:

$$V := V \cup \Delta V$$

Incremental View Maintenance

- Suppose we have computed the view:

$$V := [\text{monotone query here}]$$

- One, or several of the base relations is updated: e.g. we insert new tuples:

$$R := R \cup \Delta R$$

- Then the view needs to be incremented with some new tuples:

$$V := V \cup \Delta V$$

- Problem: compute ΔV without recomputing V

Incremental View Maintenance

Fix a USPJ query Q with inputs R_1, R_2, \dots :

$$Q(R_1, R_2, \dots)$$

The delta query ΔQ is any query that has the property:

$$Q(R_1 \cup \Delta R_1, R_2 \cup \Delta R_2, \dots) = Q(R_1, R_2, \dots) \cup \Delta Q(R_1, \Delta R_1, R_2, \Delta R_2, \dots)$$

Incremental View Maintenance

Example 1:

$V(x,y) :- R(x,z), S(z,y)$

If $R \leftarrow R \cup \Delta R$ then what is $\Delta V(x,y)$?

Incremental View Maintenance

Example 1:

$V(x,y) :- R(x,z), S(z,y)$

If $R \leftarrow R \cup \Delta R$ then what is $\Delta V(x,y)$?

$\Delta V(x,y) :- \Delta R(x,z), S(z,y)$

Incremental View Maintenance

Example 2:

$V(x,y) :- R(x,z), S(z,y)$

If $R \leftarrow R \cup \Delta R$ and $S \leftarrow S \cup \Delta S$
then what is $\Delta V(x,y)$?

Incremental View Maintenance

Example 2:

$V(x,y) :- R(x,z), S(z,y)$

If $R \leftarrow R \cup \Delta R$ and $S \leftarrow S \cup \Delta S$
then what is $\Delta V(x,y)$?

$\Delta V(x,y) :- \Delta R(x,z), S(z,y)$

$\Delta V(x,y) :- R(x,z), \Delta S(z,y)$

$\Delta V(x,y) :- \Delta R(x,z), \Delta S(z,y)$

Incremental View Maintenance

Example 3:

$V(x,y) :- T(x,z), T(z,y)$

If $T \leftarrow T \cup \Delta T$
then what is $\Delta V(x,y)$?

Incremental View Maintenance

Example 3:

$V(x,y) :- T(x,z), T(z,y)$

If $T \leftarrow T \cup \Delta T$
then what is $\Delta V(x,y)$?

$\Delta V(x,y) :- \Delta T(x,z), T(z,y)$

$\Delta V(x,y) :- T(x,z), \Delta T(z,y)$

$\Delta V(x,y) :- \Delta T(x,z), \Delta T(z,y)$

Incremental View Maintenance

Fix a USPJ query Q with inputs R_1, R_2, \dots :

$$Q(R_1, R_2, \dots)$$

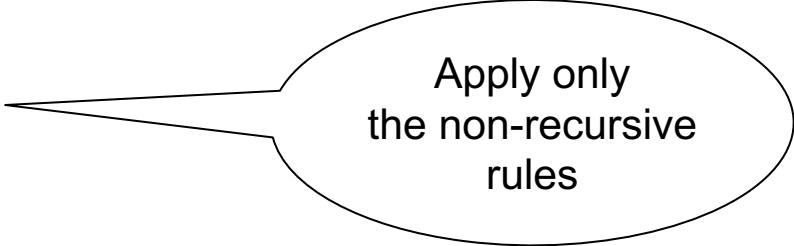
The delta query ΔQ is any query that has the property:

$$Q(R_1 \cup \Delta R_1, R_2 \cup \Delta R_2, \dots) = Q(R_1, R_2, \dots) \cup \Delta Q(R_1, \Delta R_1, R_2, \Delta R_2, \dots)$$

Semi-Naïve Evaluation Algorithm

$IDB = SPJU(\emptyset)$

$\Delta IDB = IDB$



Apply only
the non-recursive
rules

Loop

$\Delta IDB := \Delta SPJU(IDB, \Delta IDB) -- IDB$

$IDB := IDB \cup \Delta IDB$

if ($\Delta IDB = \emptyset$) then break

End Loop

From Naïve to Semi-Naive

Naïve Algorithm:

$$IDB_0 \subseteq IDB_1 \subseteq IDB_2 \subseteq \dots$$

$IDB_0 := \emptyset$

$t := 0$

Loop

$IDB_{t+1} := SPJU(IDB_t)$

if ($IDB_{t+1} = IDB_t$) then break

$t := t+1$

End Loop

From Naïve to Semi-Naive

Naïve Algorithm:

$$IDB_0 \subseteq IDB_1 \subseteq IDB_2 \subseteq \dots$$

$$\Delta IDB_t = IDB_t - IDB_{t-1}$$

$IDB_0 := \emptyset$

$t := 0$

Loop

$IDB_{t+1} := SPJU(IDB_t)$

if ($IDB_{t+1} = IDB_t$) then break

$t := t+1$

End Loop

From Naïve to Semi-Naive

Naïve Algorithm:

$$IDB_0 \subseteq IDB_1 \subseteq IDB_2 \subseteq \dots$$

$$\Delta IDB_t = IDB_t - IDB_{t-1}$$

```
IDB0 := ∅  
t := 0  
Loop  
  IDBt+1 := SPJU(IDBt)  
  if (IDBt+1 = IDBt) then break  
  t := t+1  
End Loop
```

```
IDB0 := ∅; ΔIDB0 := ∅;  
t := 0  
Loop  
  ΔIDBt+1 := SPJU(IDBt) -- IDBt  
  IDBt+1 := IDBt ∪ ΔIDBt+1  
  if (ΔIDBt+1 = ∅) then break  
  t := t+1  
End Loop
```


From Naïve to Semi-Naive

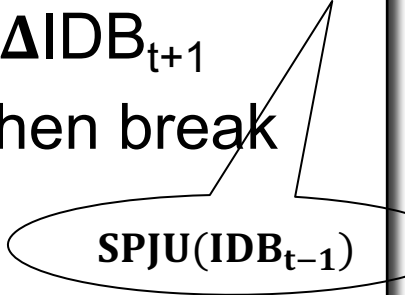
Naïve Algorithm:

$$IDB_0 \subseteq IDB_1 \subseteq IDB_2 \subseteq \dots$$

$$\Delta IDB_t = IDB_t - IDB_{t-1}$$

```
IDB0 := ∅  
t := 0  
Loop  
  IDBt+1 := SPJU(IDBt)  
  if (IDBt+1 = IDBt) then break  
  t := t+1  
End Loop
```

```
IDB0 := ∅; ΔIDB0 := ∅;  
t := 0  
Loop  
  ΔIDBt+1 := SPJU(IDBt) -- IDBt  
  IDBt+1 := IDBt ∪ ΔIDBt+1  
  if (ΔIDBt+1 = ∅) then break  
  t := t+1  
End Loop
```



From Naïve to Semi-Naive

Naïve Algorithm:

$$IDB_0 \subseteq IDB_1 \subseteq IDB_2 \subseteq \dots$$

$IDB_0 := \emptyset$

$t := 0$

Loop

$IDB_{t+1} := SPJU(IDB_t)$

if $(IDB_{t+1} = IDB_t)$ then break

$t := t+1$

End Loop

$$\Delta IDB_t = IDB_t - IDB_{t-1}$$

$IDB_0 := \emptyset; \Delta IDB_0 := \emptyset;$

$t := 0$

$SPJU(IDB_{t-1} \cup \Delta IDB_t)$

Loop

$\Delta IDB_{t+1} := SPJU(IDB_t) - IDB_t$

$IDB_{t+1} := IDB_t \cup \Delta IDB_{t+1}$

if $(\Delta IDB_{t+1} = \emptyset)$ then break

$t := t+1$

$SPJU(IDB_{t-1})$

End Loop

From Naïve to Semi-Naive

Naïve Algorithm:

$$IDB_0 \subseteq IDB_1 \subseteq IDB_2 \subseteq \dots$$

$$\Delta IDB_t = IDB_t - IDB_{t-1}$$

```
IDB0 := ∅; ΔIDB0 := ∅;  
t := 0  
Loop  
  ΔIDBt+1 := SPJU(IDBt) -- IDBt  
  IDBt+1 := IDBt ∪ ΔIDBt+1  
  if (ΔIDBt+1 = ∅) then break  
  t := t+1  
End Loop
```

$\Delta SPJU(IDB_{t-1}, \Delta IDB_t)$

$SPJU(IDB_{t-1} \cup \Delta IDB_t)$

$SPJU(IDB_{t-1})$

From Naïve to Semi-Naive

Naïve Algorithm:

$$IDB_0 \subseteq IDB_1 \subseteq IDB_2 \subseteq \dots$$

$$\Delta IDB_t = IDB_t - IDB_{t-1}$$

```
IDB0 := ∅; ΔIDB0 := ∅;  
t := 0  
Loop  
  ΔIDBt+1 := SPJU(IDBt) -- IDBt  
  IDBt+1 := IDBt ∪ ΔIDBt+1  
  if (ΔIDBt+1 = ∅) then break  
  t := t+1  
End Loop
```

$\Delta SPJU(IDB_{t-1}, \Delta IDB_t)$

$SPJU(IDB_{t-1} \cup \Delta IDB_t)$

$SPJU(IDB_{t-1})$

But need to start from $t=1$, since IDB_{-1} undefined

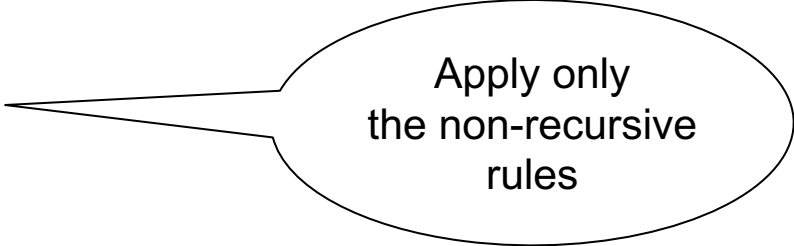
Semi-Naïve Algorithm

```
IDB1 := SPJU(∅); ΔIDB1 := IDB1;  
t := 1  
Loop  
  ΔIDBt+1 := ΔSPJU(IDBt-1, ΔIDBt) -- IDBt  
  IDBt+1 := IDBt ∪ ΔIDBt+1  
  if (ΔIDBt+1 = ∅) then break  
  t := t+1  
End Loop
```

Semi-Naïve Evaluation Algorithm

$IDB = SPJU(\emptyset)$

$\Delta IDB = IDB$



Apply only
the non-recursive
rules

Loop

$\Delta IDB := \Delta SPJU(IDB, \Delta IDB) -- IDB$

$IDB := IDB \cup \Delta IDB$

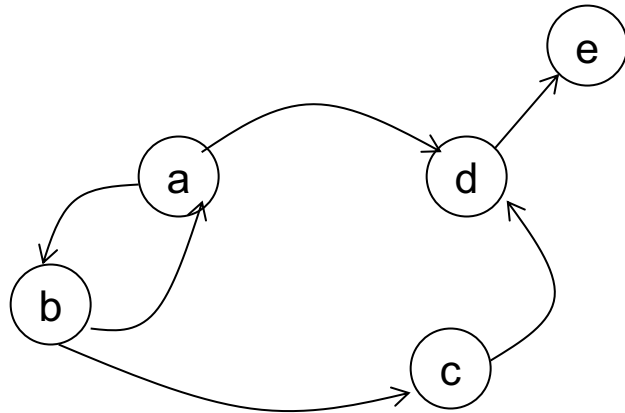
if ($\Delta IDB = \emptyset$) then break

End Loop

$T(x,y) :- R(x,y)$

$T(x,y) :- R(x,z), T(z,y)$

Example: Linear Recursion



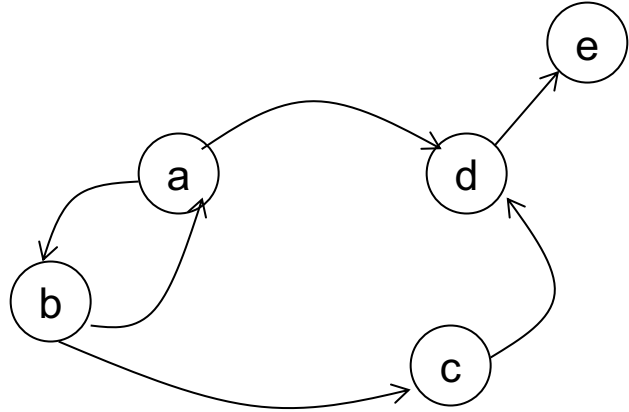
R=

a	b
b	a
b	c
a	d
c	d
d	e

$T(x,y) :- R(x,y)$

$T(x,y) :- R(x,z), T(z,y)$

Example



R=

a	b
b	a
b	c
a	d
c	d
d	e

$T := R \quad \Delta T := T$

Loop

$\Delta T(x,y) := \Pi_{x,y}(R(x,z) \bowtie \Delta T(z,y)) - T(x,y)$

$T := T \cup \Delta T$

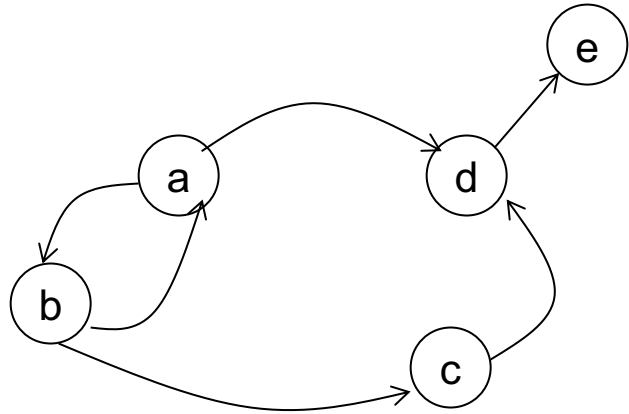
if $(\Delta T = \emptyset)$ then break

End Loop

$T(x,y) :- R(x,y)$

$T(x,y) :- R(x,z), T(z,y)$

Example



$T := R \quad \Delta T := T$

Loop

$\Delta T(x,y) := \Pi_{x,y}(R(x,z) \bowtie \Delta T(z,y)) - T(x,y)$

$T := T \cup \Delta T$

if $(\Delta T = \emptyset)$ then break

End Loop

R=

a	b
b	a
b	c
a	d
c	d
d	e

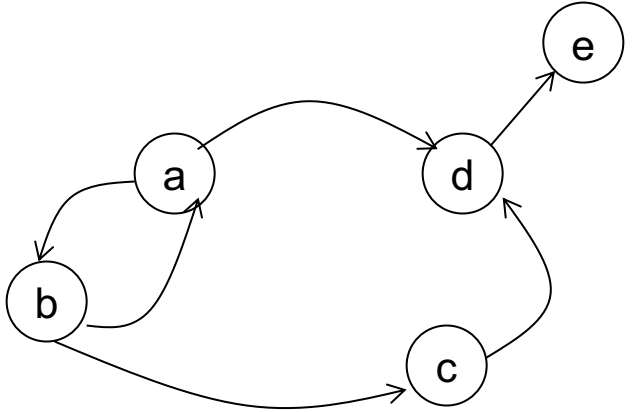
Step 0:

ΔT		T	
a	b	a	b
b	a	b	a
b	c	b	c
a	d	a	d
c	d	c	d
d	e	d	e

$T(x,y) :- R(x,y)$
 $T(x,y) :- R(x,z), T(z,y)$

$T := R \quad \Delta T := T$
 Loop
 $\Delta T(x,y) := \Pi_{x,y}(R(x,z) \bowtie \Delta T(z,y)) - T(x,y)$
 $T := T \cup \Delta T$
 if $(\Delta T = \emptyset)$ then break
 End Loop

Example



R=

a	b
b	a
b	c
a	d
c	d
d	e

Step 0:

ΔT		T	
a	b	a	b
b	a	b	a
b	c	b	c
a	d	a	d
c	d	c	d
d	e	d	e

Step 1:

ΔT		T	
a	a	a	b
b	b	b	a
a	c	b	c
b	d	a	d
a	e	c	d
c	e	d	e

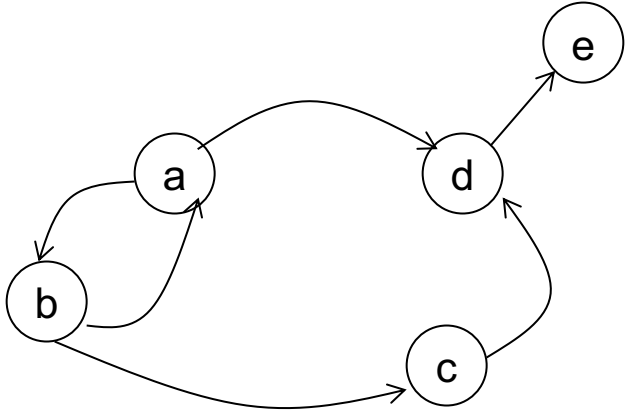
T

a	b
b	a
b	c
a	d
c	d
d	e
a	a
b	b
a	c
b	d
a	e
c	e

$T(x,y) :- R(x,y)$
 $T(x,y) :- R(x,z), T(z,y)$

$T := R \quad \Delta T := T$
 Loop
 $\Delta T(x,y) := \Pi_{x,y}(R(x,z) \bowtie \Delta T(z,y)) - T(x,y)$
 $T := T \cup \Delta T$
 if $(\Delta T = \emptyset)$ then break
 End Loop

Example



R=

a	b
b	a
b	c
a	d
c	d
d	e

Step 0:

ΔT		T	
a	b	a	b
b	a	b	a
b	c	b	c
a	d	a	d
c	d	c	d
d	e	d	e

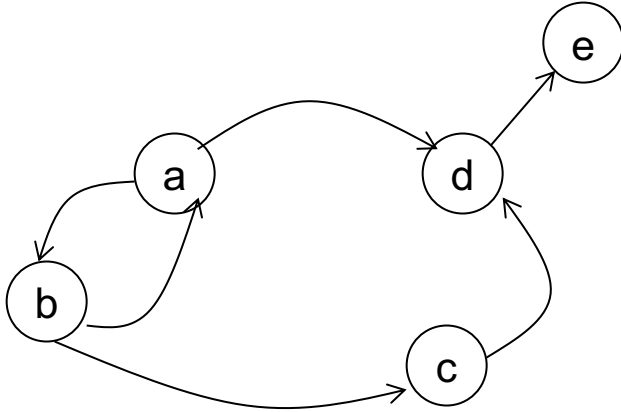
Step 1:

ΔT		T	
a	a	a	b
b	b	b	a
a	c	b	c
b	d	a	d
a	e	c	d
c	e	a	e
		c	e

Step 2:

ΔT		T	
d	e	a	b
		b	a
		b	c
		a	d
		c	d
		d	e
		a	a
		b	b
		a	c
		b	d
		a	e
		c	e
		d	e

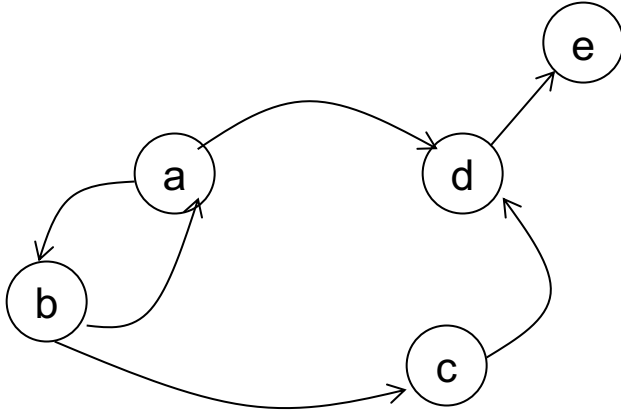
Example: Non-linear Recursion



$T(x,y) :- R(x,y)$

$T(x,y) :- T(x,z), T(z,y)$

Example: Non-linear Recursion



$T(x,y) \text{ :- } R(x,y)$

$T(x,y) \text{ :- } T(x,z), T(z,y)$

Semi-naïve algorithm:

$T \text{ := } R \quad \Delta T \text{ := } T$

Loop

$\Delta T(x,y) \text{ := } (\Pi_{x,y}(T(x,z) \bowtie \Delta T(z,y)) \cup \Pi_{x,y}(\Delta T(x,z) \bowtie T(z,y)) \cup \Pi_{x,y}(\Delta T(x,z) \bowtie \Delta T(z,y))) - T(x,y)$

$T \text{ := } T \cup \Delta T$

if $(\Delta T = \emptyset)$ then break

End Loop

Example: Same Generation

$SG(x,y) :- \text{ParentChild}(p,x), \text{ParentChild}(p,y)$

$SG(x,y) :- \text{ParentChild}(p,x), SG(p,q), \text{ParentChild}(q,y)$

Example: Same Generation

$SG(x,y) :- \text{ParentChild}(p,x), \text{ParentChild}(p,y)$

$SG(x,y) :- \text{ParentChild}(p,x), SG(p,q), \text{ParentChild}(q,y)$

$SG := \text{ParentChild} \quad \Delta SG := SG$

Loop

$\Delta SG(x,y) := \Pi_{x,y}(\text{ParentChild}(p,x) \bowtie \Delta SG(p,q) \bowtie \text{ParentChild}(q,y)) - SG(x,y)$

$SG := SG \cup \Delta SG$

if $(\Delta SG = \emptyset)$ then break

End Loop

Discussion

- Most datalog engines implement the semi-naïve algorithm
- Notice: it only works when the recursion has a monotone body

Summary

- Datalog = light-weight syntax, recursion
- Several optimizations
- Limitations:
 - Monotone queries work great
 - Non-monotone queries: stratification
- SQL: supports limited recursion only