

# CSE544: Principles of Database Systems

Part III: Database Theory  
Datalog

# Why DB Theory

- Discuss the urgency of parallelism in the paper
- Discuss why data-centric approach to parallel computing
- Theory is key for understanding this approach

# Outline of DB Theory

- Datalog – this lecture
- Query complexity – next week
- Static analysis (query equivalence)
- Advanced optimizations (semijoin reduction)

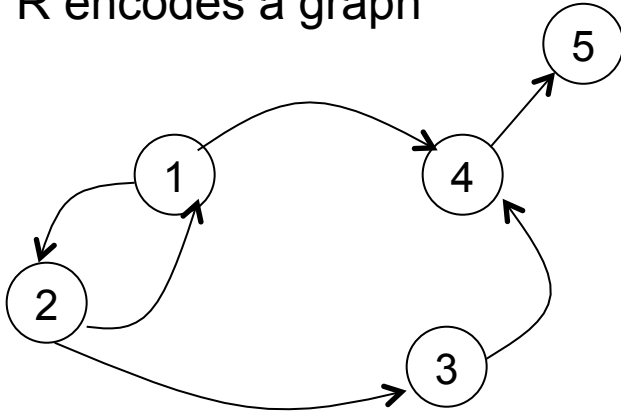
# Datalog

Review the following basic concepts from Lecture 2:

- Fact
- Rule
- Head and body of a rule
- Existential variable
- Head variable

# Simple datalog programs

R encodes a graph



R=

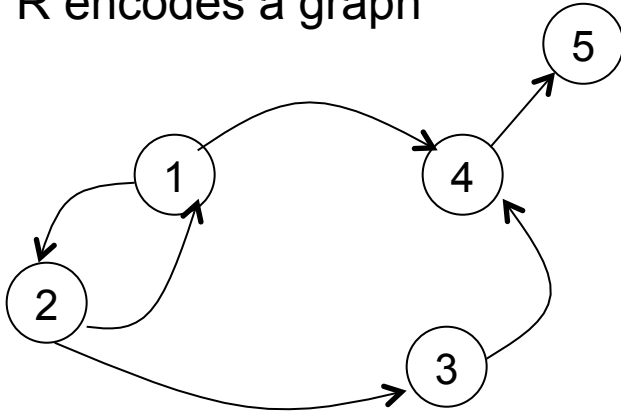
1	2
2	1
2	3
1	4
3	4
4	5

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

What does it compute?

# Simple datalog programs

R encodes a graph



R=

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

Initially:  
T is empty.

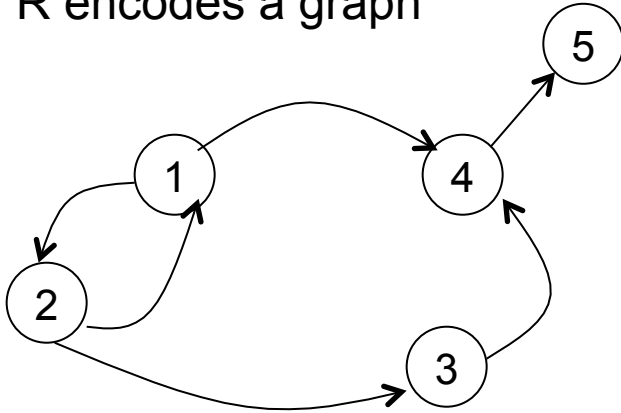


$T(x,y) :- R(x,y)$   
 $T(x,y) :- R(x,z), T(z,y)$

What does  
it compute?

# Simple datalog programs

R encodes a graph



R =

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

Initially:  
T is empty.



First iteration:  
T =

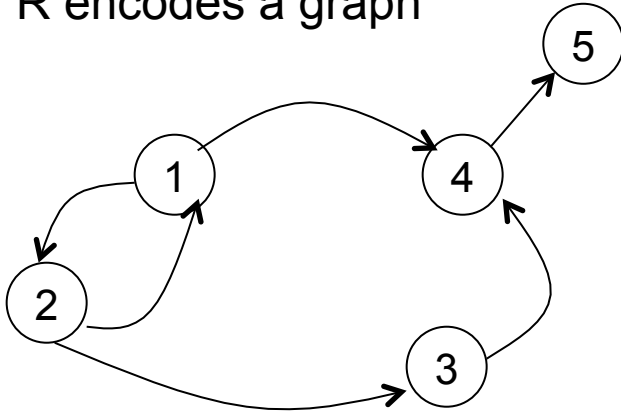
1	2
2	1
2	3
1	4
3	4
4	5

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

What does  
it compute?

# Simple datalog programs

R encodes a graph



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

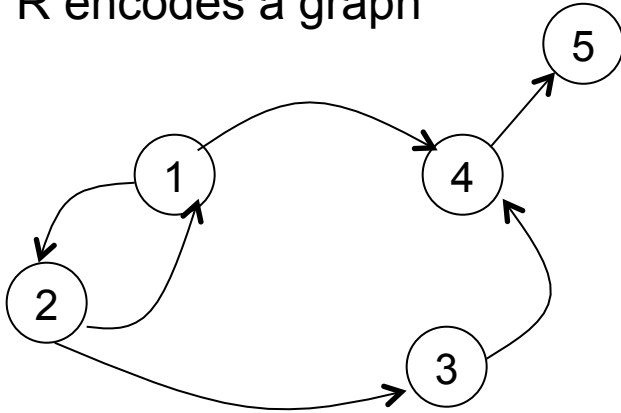
$T(x,y) \text{ :- } R(x,y)$   
 $T(x,y) \text{ :- } R(x,z), T(z,y)$

What does it compute?



# Simple datalog programs

R encodes a graph



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

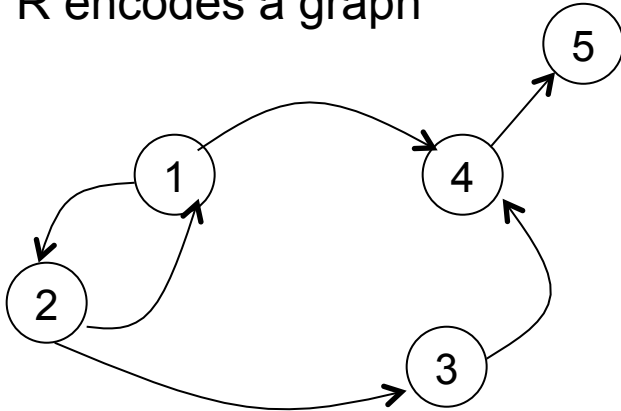
$T(x,y) \text{ :- } R(x,y)$   
 $T(x,y) \text{ :- } R(x,z), T(z,y)$

What does  
it compute?

Done

# Simple datalog programs

R encodes a graph



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

Discovered  
3 times!

Discovered  
twice

Done

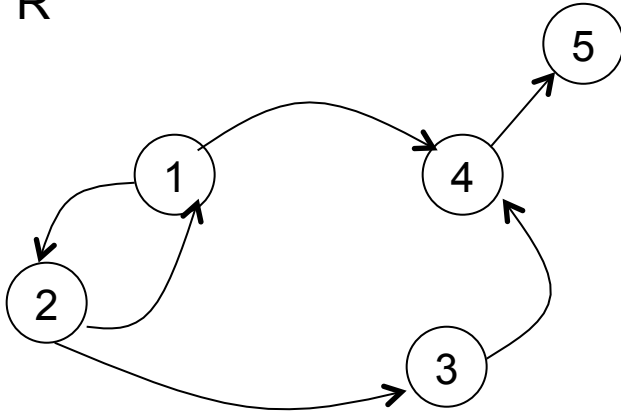
$$T(x,y) \text{ :- } R(x,y)$$

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

What does  
it compute?

# Simple datalog programs

R



R=

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

Alternative ways to compute TC:

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

Right linear

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

Left linear

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

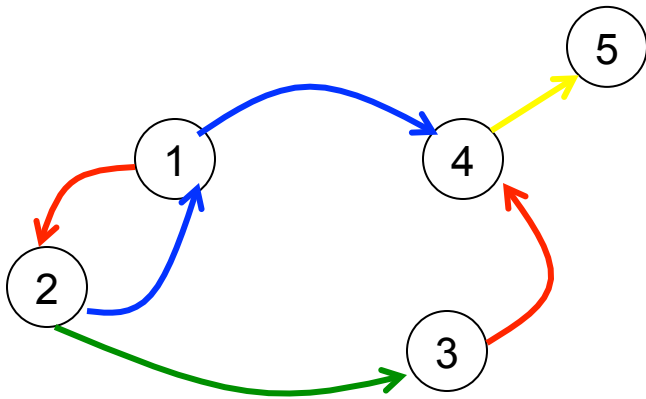
Non-linear

Discuss pros/cons in class

# Simple datalog programs

Compute TC (ignoring color):

R encodes a **colored** graph



Compute pairs of nodes connected by the same color (e.g. (2,4))

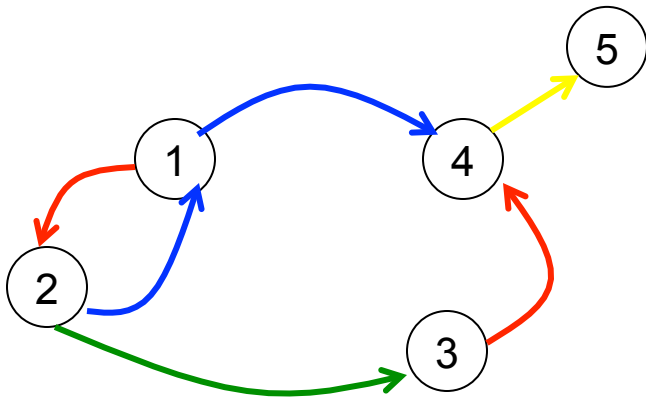
R=

1	Red	2
2	Blue	1
2	Green	3
1	Blue	4
3	Red	4
4	Yellow	5

# Simple datalog programs

Compute TC (ignoring color):

R encodes a **colored** graph



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

Compute pairs of nodes connected by the same color (e.g. (2,4))

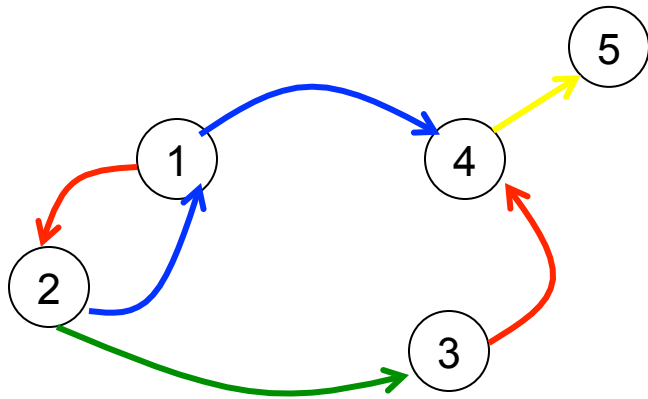
R=

1	Red	2
2	Blue	1
2	Green	3
1	Blue	4
3	Red	4
4	Yellow	5

# Simple datalog programs

Compute TC (ignoring color):

R encodes a **colored** graph



R=

1	Red	2
2	Blue	1
2	Green	3
1	Blue	4
3	Red	4
4	Yellow	5

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

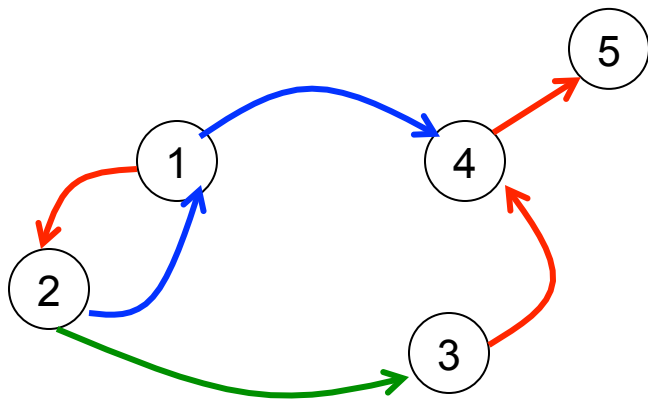
Compute pairs of nodes connected by the same color (e.g. (2,4))

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

# Simple datalog programs

R, G, B encodes a 3-colored graph

What does this program compute in general?



R=

1	2
3	4
4	5

G=

2	3
---	---

B=

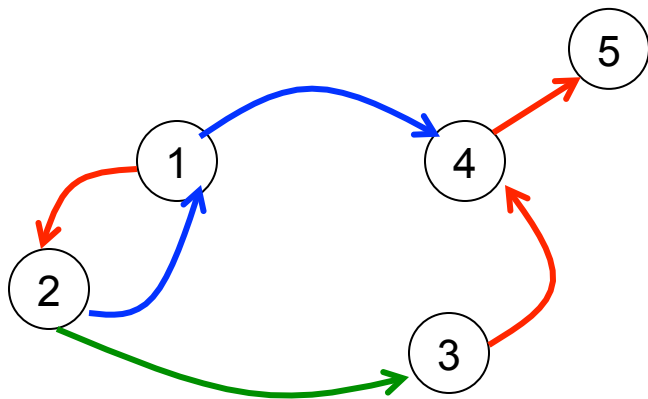
2	1
1	4

```
T(x,y) :- R(x,y)
T(x,y) :- G(x,y)
S(x,y) :- T(x,z), B(z,y)
S(x,y) :- S(x,z), B(z,y)
T(x,y) :- S(x,z), R(z,y)
T(x,y) :- S(x,z), G(z,y)
Answer(x,y) :- S(x,y)
```

# Simple datalog programs

R, G, B encodes a 3-colored graph

What does this program compute in general?



R=

1	2
3	4
4	5

G=

2	3
---	---

B=

2	1
1	4

```
T(x,y) :- R(x,y)
T(x,y) :- G(x,y)
S(x,y) :- T(x,z), B(z,y)
S(x,y) :- S(x,z), B(z,y)
T(x,y) :- S(x,z), R(z,y)
T(x,y) :- S(x,z), G(z,y)
Answer(x,y) :- S(x,y)
```

Answer: it computes pairs of nodes connected by a path spelling out certain regular expressions:

- $S = ((R \text{ or } G).B^+)^*$
- $T = ((R \text{ or } G).B^+)^*.(R \text{ or } G)$



# Syntax of Datalog Programs

The schema consists of two sets of relations:

- Extensional Database (EDB):  $R_1, R_2, \dots$
- Intentional Database (IDB):  $P_1, P_2, \dots$

A datalog program  $\mathbf{P}$  has the form:

$\mathbf{P}$ :

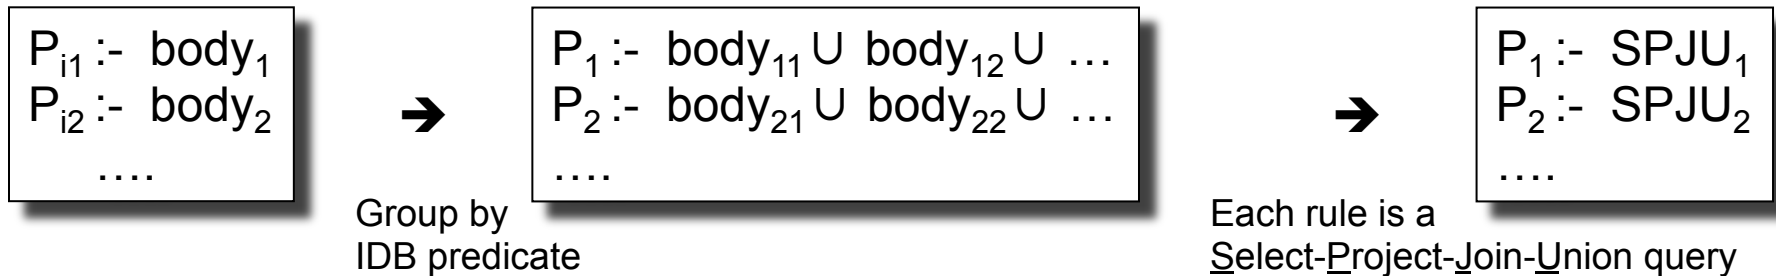
$P_{i1}(x_{11}, x_{12}, \dots) :- \text{body}_1$
$P_{i2}(x_{21}, x_{22}, \dots) :- \text{body}_2$
.....

- Each head predicate  $P_i$  is an IDB
- Each body is a conjunction of IDB and/or EDB predicates
- See lecture 2

Note: no negation (yet)! Recursion OK.

# Naïve Datalog Evaluation Algorithm

Datalog program:



Naïve datalog evaluation algorithm:

```

P1 = P2 = ... = ∅
Loop
  NewP1 = SPJU1; NewP2 = SPJU2; ...
  if (NewP1 = P1 and NewP2 = P2 and ...)
    then break
  P1 = NewP1; P2 = NewP2; ...
Endloop
    
```

Example:

```

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

→  $T(x,y) \text{ :- } R(x,y) \cup \Pi_{xy}(R(x,z) \bowtie T(z,y))$

```

T = ∅
Loop
  NewT(x,y) = R(x,y) ∪ Πxy(R(x,z) ⋈ T(z,y))
  if (NewT = T)
    then break
  T = NewT
Endloop
    
```

# Problem with the Naïve Algorithm

- The same facts are discovered over and over again
- The semi-naïve algorithm tries to reduce the number of facts discovered multiple times

# Background: Incremental View Maintenance

Let  $V$  be a view computed by one datalog rule (no recursion)

$V \text{ :- body}$

If (some of) the relations are updated:  $R_1 \leftarrow R_1 \cup \Delta R_1, R_2 \leftarrow R_2 \cup \Delta R_2, \dots$

Then the view is also modified as follows:  $V \leftarrow V \cup \Delta V$

**Incremental view maintenance:** compute  $\Delta V$  (without having to recompute  $V$ )

# Background: Incremental View Maintenance

Let  $V$  be a view computed by one datalog rule (no recursion)

$V :- \text{body}$

If (some of) the relations are updated:  $R_1 \leftarrow R_1 \cup \Delta R_1, R_2 \leftarrow R_2 \cup \Delta R_2, \dots$

Then the view is also modified as follows:  $V \leftarrow V \cup \Delta V$

**Incremental view maintenance:** compute  $\Delta V$  (without having to recompute  $V$ )

Solution: by examples...

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

$\Delta V(x,y) :- ???$

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

$\Delta W(x,y) :- ???$

$W(x,y) :- R(x,u), S(u,v), T(v,y)$

.....

# Background: Incremental View Maintenance

Let  $V$  be a view computed by one datalog rule (no recursion)

$V :- \text{body}$

If (some of) the relations are updated:  $R_1 \leftarrow R_1 \cup \Delta R_1, R_2 \leftarrow R_2 \cup \Delta R_2, \dots$

Then the view is also modified as follows:  $V \leftarrow V \cup \Delta V$

**Incremental view maintenance:** compute  $\Delta V$  (without having to recompute  $V$ )

Solution: by examples...

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

$\Delta V(x,y) :- ???$

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

$\Delta W(x,y) :- ???$

$W(x,y) :- R(x,u), S(u,v), T(v,y)$

.....

# Background: Incremental View Maintenance

Let  $V$  be a view computed by one datalog rule (no recursion)

$V \text{ :- body}$

If (some of) the relations are updated:  $R_1 \leftarrow R_1 \cup \Delta R_1, R_2 \leftarrow R_2 \cup \Delta R_2, \dots$

Then the view is also modified as follows:  $V \leftarrow V \cup \Delta V$

**Incremental view maintenance:** compute  $\Delta V$  (without having to recompute  $V$ )

Solution: by examples...

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

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

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

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

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

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

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

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

$W(x,y) \text{ :- } R(x,u), S(u,v), T(v,y)$

....

Note: one rule may generate multiple  $\Delta$ -rules

# Semi-naïve Evaluation Algorithm

Separate the Datalog program into the non-recursive, and the recursive part.  
Each  $P_i$  defined by non-recursive-SPJU $_i$  and (recursive-)SPJU $_i$ .

```
P1 = P2 = ... = ∅,  
ΔP1 = non-recursive-SPJU1, ΔP2 = non-recursive-SPJU2, ...  
Loop  
  ΔP1 = Δ SPJU1; ΔP2 = ΔSPJU2; ...  
  if (ΔP1 = ∅ and ΔP2 = ∅ and ...)  
    then break  
  P1 = P1 ∪ ΔP1; P2 = P2 ∪ ΔP2; ...  
Endloop
```

Example:

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

```
T = ∅, ΔT = R  
Loop  
  ΔT(x,y) = Πxy(R(x,z) ⋈ ΔT(z,y))  
  if (ΔT = ∅)  
    then break  
  T = T ∪ ΔT  
Endloop
```

Note: for any linear datalog programs,  
the semi-naïve algorithm has only  
one  $\Delta$ -rule for each rule!



# Discussion in Class

How would you compute the transitive closure of a very large graph  $R(x,y)$ ?

- Assume a single server
- Assume a shared nothing architecture

Right linear TC

$$\begin{aligned} T(x,y) &:- R(x,y) \\ T(x,y) &:- R(x,z), T(z,y) \end{aligned}$$

Non-linear TC

$$\begin{aligned} T(x,y) &:- R(x,y) \\ T(x,y) &:- T(x,z), T(z,y) \end{aligned}$$

# Discussion in Class

The *Declarative Imperative* paper:

- What are the extensions to datalog in Dedalus?
- What is the main usage of Dedalus described in the paper? Discuss some applications, discuss what's missing.

# Semantics of a Datalog Program

Three different, equivalent semantics:

- Minimal model semantics
- Least fixpoint semantics
- Proof-theoretic semantics

# Minimal Model Semantics (1/2)

To each rule r:  $P(x_1 \dots x_k) :- R_1(\dots), R_2(\dots), \dots$

All variables in the rule

Associate the logical sentence  $\Sigma_r$ :  $\forall z_1 \dots \forall z_n. [(R_1(\dots) \wedge R_2(\dots) \wedge \dots) \rightarrow P(\dots)]$

Same as:  $\forall x_1 \dots \forall x_k. [\exists y_1 \dots \exists y_m. (R_1(\dots) \wedge R_2(\dots) \wedge \dots) \rightarrow P(\dots)]$

Head variables

Existential variables

**Definition.** If  $P$  is a datalog program,  $\Sigma_P$  is the set of all logical sentences associated to its rules.

Example. Rule:  $T(x,y) :- R(x,z), T(z,y)$

Sentence:  $\forall x. \forall y. \forall z. (R(x,z) \wedge T(z,y) \rightarrow T(x,y))$   
 $\equiv \forall x. \forall y. (\exists z. R(x,z) \wedge T(z,y) \rightarrow T(x,y))$

# Minimal Model Semantics (1/2)

**Definition.** A pair  $(I, J)$  where  $I$  is an EDB and  $J$  is an IDB is a *model* for  $P$ , if  $(I, J) \models \Sigma_P$

**Definition.** Given an EDB database instance  $I$  and a datalog program  $P$ , the minimal model, denoted  $J = \mathbf{P}(I)$  is a minimal database instance  $J$  s.t.  $(I, J) \models \Sigma_P$

**Theorem.** The minimal model always exists, and is unique.

# Minimal Model Semantics (1/2)

**Definition.** A pair  $(I,J)$  where  $I$  is an EDB and  $J$  is an IDB is a *model* for  $P$ , if  $(I,J) \models \Sigma_P$

**Definition.** Given an EDB database instance  $I$  and a datalog program  $P$ , the minimal model, denoted  $J = \mathbf{P}(I)$  is a minimal database instance  $J$  s.t.  $(I,J) \models \Sigma_P$

**Theorem.** The minimal model always exists, and is unique.

Example:



Which of these IDBs are *models*?  
Which are *minimal models*?

R=

1	2
2	3
3	4
4	5

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

T=

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

T=

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

T=

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

All pairs of distinct nodes

# Minimal Fixpoint Semantics (1/2)

**Definition.** Fix an EDB  $I$ , and a datalog program  $P$ .

The *immediate consequence* operator  $T_P$  is defined as follows.

For any IDB  $J$ :

$T_P(J)$  = all IDB facts that are immediate consequences from  $I$  and  $J$ .

**Fact.** For any datalog program  $P$ , the immediate consequence operator is monotone. In other words, if  $J_1 \subseteq J_2$  then  $T_P(J_1) \subseteq T_P(J_2)$ .

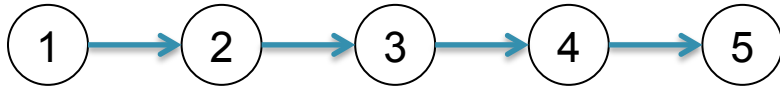
**Theorem.** The immediate consequence operator has a unique, minimal fixpoint  $J$ :  $\text{fix}(T_P) = J$ , where  $J$  is the minimal instance with the property  $T_P(J) = J$ .

Proof: using Knaster-Tarski's theorem for monotone functions.

The fixpoint is given by:

$\text{fix}(T_P) = J_0 \cup J_1 \cup J_2 \cup \dots$  where  $J_0 = \emptyset$ ,  $J_{k+1} = T_P(J_k)$

# Minimal Fixpoint Semantics (2/2)



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

R=

1	2
2	3
3	4
4	5

T =

--	--

$J_0 = \emptyset$

$J_1 = T_P(J_0)$

1	2
2	3
3	4
4	5

$J_2 = T_P(J_1)$

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

$J_3 = T_P(J_2)$

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

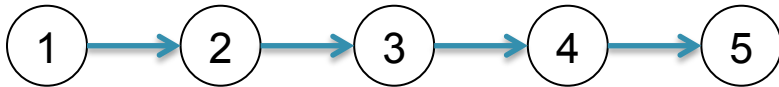
$J_4 = T_P(J_3)$

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



# Proof Theoretic Semantics

Every fact in the IDB has a *derivation tree*, or *proof tree* justifying its existence.

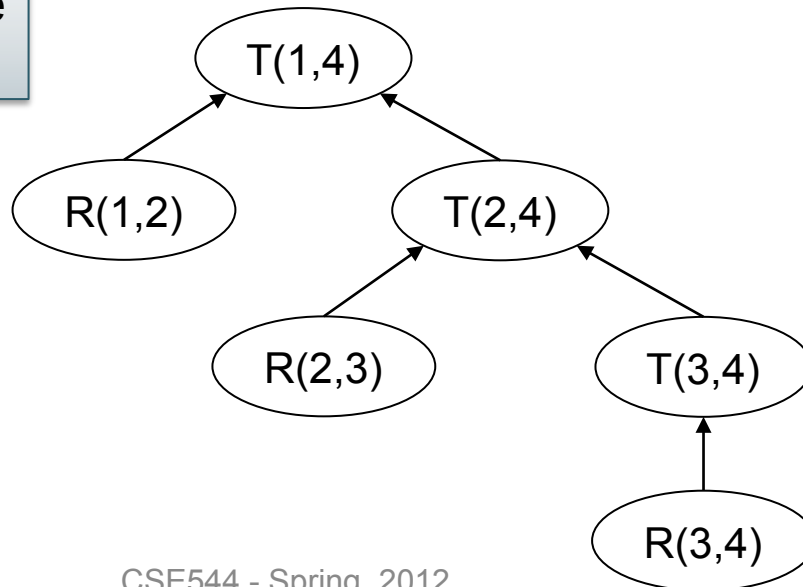


$T(x,y) \text{ :- } R(x,y)$   
 $T(x,y) \text{ :- } R(x,z), T(z,y)$

R=

1	2
2	3
3	4
4	5

Derivation tree  
of  $T(1,4)$



# Adding Negation: Datalog<sup>-</sup>

**Example:** compute the complement of the transitive closure

```
T(x,y) :- R(x,y)
T(x,y) :- T(x,z), R(z,y)
CT(x,y) :- Node(x), Node(y), not T(x,y)
```

What does this mean??

# Recursion and Negation Don't Like Each Other

EDB:  $I = \{ R(a) \}$

$S(x) :- R(x), \text{ not } T(x)$   
 $T(x) :- R(x), \text{ not } S(x)$

Which IDBs are models of **P**?

$J_1 = \{ \}$

$J_2 = \{ S(a) \}$

$J_3 = \{ T(a) \}$

$J_4 = \{ S(a), T(a) \}$

# Recursion and Negation Don't Like Each Other

EDB:  $I = \{ R(a) \}$

$S(x) :- R(x), \text{ not } T(x)$   
 $T(x) :- R(x), \text{ not } S(x)$

Which IDBs are models of **P**?

$J_1 = \{ \}$

No: both  
rules fail

$J_2 = \{ S(a) \}$

Yes: the facts in  $J_2$  are  
 $R(a), S(a), \neg T(a)$   
and both rules are *true*.

$J_3 = \{ T(a) \}$

Yes

$J_4 = \{ S(a), T(a) \}$

Yes

There is no minimal model!

# Recursion and Negation Don't Like Each Other

EDB:  $I = \{ R(a) \}$

$S(x) :- R(x), \text{ not } T(x)$   
 $T(x) :- R(x), \text{ not } S(x)$

Which IDBs are models of **P**?

$J_1 = \{ \}$

No: both  
rules fail

$J_2 = \{ S(a) \}$

Yes: the facts in  $J_2$  are  
 $R(a), S(a), \neg T(a)$   
and both rules are *true*.

$J_3 = \{ T(a) \}$

Yes

$J_4 = \{ S(a), T(a) \}$

Yes

There is no minimal model!

There is no minimal fixpoint!  
(Why does Knaster-Tarski's  
theorem fail?)

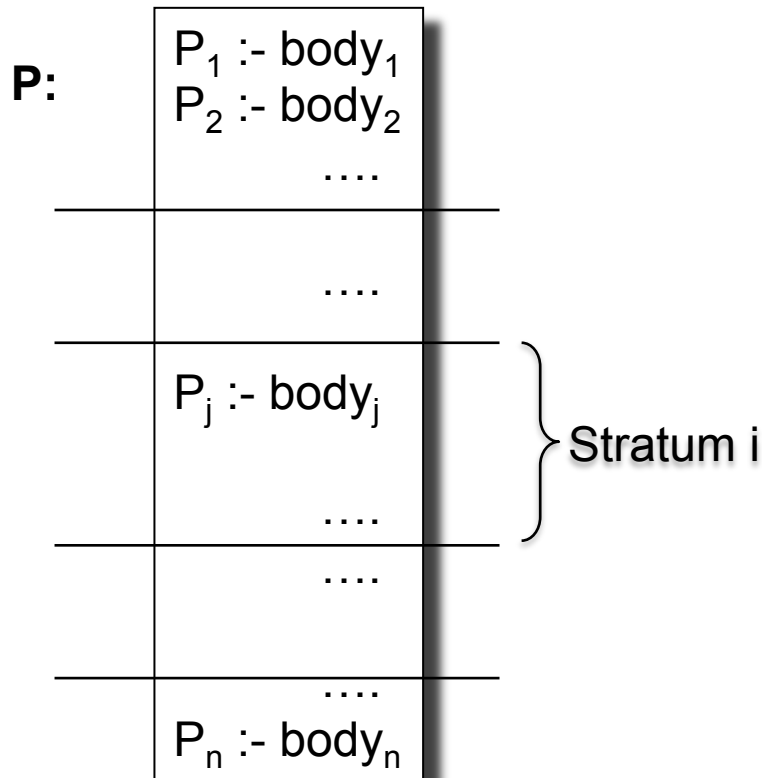
# Adding Negation: datalog<sup>-</sup>

- **Solution 1: Stratified Datalog<sup>-</sup>**
  - Insist that the program be stratified: rules are partitioned into strata, and an IDB predicate that occurs only in strata  $\leq k$  may be negated in strata  $\geq k+1$
- **Solution 2: Inflationary-fixpoint Datalog<sup>-</sup>**
  - Compute the fixpoint of  $J \cup T_P(J)$
  - Always terminates (why ?)
- **Solution 3: Partial-fixpoint Datalog<sup>-,\*</sup>**
  - Compute the fixpoint of  $T_P(J)$
  - May not terminate

# Stratified datalog<sup>-</sup>

A datalog<sup>-</sup> program is *stratified* if its rules can be partitioned into k strata, such that:

- If an IDB predicate P appears negated in a rule in stratum i, then it can only appear in the head of a rule in strata 1, 2, ..., i-1



Note: a datalog<sup>-</sup> program either is stratified or it ain't!

Which programs are stratified?

$T(x,y) :- R(x,y)$   
 $T(x,y) :- T(x,z), R(z,y)$   
 $CT(x,y) :- \text{Node}(x), \text{Node}(y), \text{not } T(x,y)$

$S(x) :- R(x), \text{not } T(x)$   
 $T(x) :- R(x), \text{not } S(x)$

# Stratified datalog<sup>-</sup>

- Evaluation algorithm for stratified datalog<sup>-</sup>:
- For each stratum  $i = 1, 2, \dots$ , do:
  - Treat all IDB's defined in prior strata as EBS
  - Evaluate the IDB's defined in stratum  $i$ , using either the naïve or the semi-naïve algorithm

Does this compute a minimal model?

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

```
CT(x,y) :- Node(x), Node(y), not T(x,y)
```



# Stratified datalog<sup>-</sup>

- Evaluation algorithm for stratified datalog<sup>-</sup>:
- For each stratum  $i = 1, 2, \dots$ , do:
  - Treat all IDB's defined in prior strata as EBS
  - Evaluate the IDB's defined in stratum  $i$ , using either the naïve or the semi-naïve algorithm

Does this compute a minimal model?

NO:

$J_1 = \{ T = \text{transitive closure, CT} = \text{its complement} \}$

$J_2 = \{ T = \text{all pairs of nodes, CT} = \text{empty} \}$

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

$CT(x,y) :- \text{Node}(x), \text{Node}(y), \text{not } T(x,y)$

# Inflationary-fixpoint datalog<sup>-</sup>

Let  $\mathbf{P}$  be any datalog<sup>-</sup> program, and  $I$  an EDB.

Let  $T_{\mathbf{P}}(J)$  be the *immediate consequence* operator.

Let  $F(J) = J \cup T_{\mathbf{P}}(J)$  be the *inflationary immediate consequence* operator.

Define the sequence:  $J_0 = \emptyset$ ,  $J_{n+1} = F(J_n)$ , for  $n \geq 0$ .

**Definition.** The inflationary fixpoint semantics of  $\mathbf{P}$  is  $J = J_n$  where  $n$  is such that  $J_{n+1} = J_n$

Why does there always exist an  $n$  such that  $J_n = F(J_n)$ ?

Find the inflationary semantics for:

```
T(x,y) :- R(x,y)
T(x,y) :- T(x,z), R(z,y)
CT(x,y) :- Node(x), Node(y), not T(x,y)
```

```
S(x) :- R(x), not T(x)
T(x) :- R(x), not S(x)
```

# Inflationary-fixpoint datalog<sup>-</sup>

- Evaluation for Inflationary-fixpoint datalog<sup>-</sup>
- Use the naïve, or the semi-naïve algorithm
- Inhibit any optimization that rely on monotonicity (e.g. out of order execution)

# Partial-fixpoint datalog<sup>¬,\*</sup>

Let  $\mathbf{P}$  be any datalog<sup>¬</sup> program, and  $I$  an EDB.

Let  $T_{\mathbf{P}}(J)$  be the *immediate consequence* operator.

Define the sequence:  $J_0 = \emptyset$ ,  $J_{n+1} = T_{\mathbf{P}}(J_n)$ , for  $n \geq 0$ .

**Definition.** The partial fixpoint semantics of  $\mathbf{P}$  is  $J = J_n$  where  $n$  is such that  $J_{n+1} = J_n$ , if such an  $n$  exists, undefined otherwise.

Find the partial fixpoint semantics for:

Note: there may not exist an  $n$  such that  $J_n = F(J_n)$

```
T(x,y) :- R(x,y)
T(x,y) :- T(x,z), R(z,y)
CT(x,y) :- Node(x), Node(y), not T(x,y)
```

```
S(x) :- R(x), not T(x)
T(x) :- R(x), not S(x)
```

# Discussion

- Which semantics does Daedalus adopt?

# Discussion

## Comparing datalog<sup>-</sup>

- Compute the complement of the transitive closure in inflationary datalog<sup>-</sup>
- Compare the expressive power of:
  - Stratified datalog<sup>-</sup>
  - Inflationary fixpoint datalog<sup>-</sup>
  - Partial fixpoint datalog<sup>-</sup>

# Discussion

## Comparing datalog<sup>-</sup>

- Compute the complement of the transitive closure in inflationary datalog<sup>-</sup>
- Compare the expressive power of stratified datalog<sup>-</sup> and inflationary datalog<sup>-</sup>

You will answer both these questions in HW3!