

# Introduction to Database Systems CSE 414

## Lecture 10: More Datalog

CSE 414 - Autumn 2018

1

## Announcements

- HW 3 due Friday
  - Upload data with DataGrip editor – see message board
  - Azure timeout for question 5:
    - Try DataGrip or SQLite
  - Remember 2 late-day policy
- Gradiance web quizzes were offline:  
WQ3 due date extended one day

CSE 414 - Autumn 2018

2

## What is Datalog?

- Another query language for relational model
  - Designed in the 80's
  - Simple, concise, elegant
  - Extends relational queries with *recursion*
- Today is a hot topic:
  - Souffle (we will use in HW4)
  - Eve <http://witheve.com/>
  - Differential datalog  
<https://github.com/frankmcsherry/differential-dataflow>
  - Beyond databases in many research projects:  
network protocols, static program analysis

CSE 414 - Autumn 2018

3



- Open-source implementation of Datalog DBMS
- Under active development
- Commercial implementations are available
  - More difficult to set up and use
- “sqlite” of Datalog
  - Set-based rather than bag-based
- Install in your VM
  - Run `sudo yum install souffle` in terminal
  - More details in upcoming HW4

CSE 414 - Autumn 2018


4

## Why bother with yet another relational query language?

CSE 414 - Autumn 2018

5

## Example: storing FB friends



Or

Person1	Person2	is_friend
Peter	John	1
John	Mary	0
Mary	Phil	1
Phil	Peter	1
...	...	...

As a graph

As a relation

We will learn the tradeoffs of different data models later this quarter

CSE 414 - Autumn 2018

6

## Compute your friends graph

p1	p2	isFriend
Peter	John	1
John	Mary	0
Mary	Phil	1
Phil	Peter	1
...	...	...

Friends(p1, p2, isFriend)

```
SELECT f.p2
FROM Friends as f
WHERE f.p1 = 'me' AND f.isFriend = 1
```

My own friends

```
SELECT f1.p2
FROM Friends as f1,
      (SELECT f.p2
       FROM Friends as f
        WHERE f.p1 = 'me' AND
              f.isFriend = 1) as f2
WHERE f1.p1 = f2.p2 AND
      f1.isFriend = 1
```

My FoF

My FoFoF... My FoFoFoF...

Datalog allows us to write recursive queries easily

CSE 414 - Autumn 2018 7

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

## Datalog: Facts and Rules

Facts = tuples in the database      Rules = queries

CSE 414 - Autumn 2018 8

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

## Datalog: Facts and Rules

Facts = tuples in the database      Rules = queries

```
.decl Actor(id:number, fname:symbol, lname:symbol)
.decl Casts(id:number, mid:number)
.decl Movie(id:number, name:symbol, year:number)

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).
```

Table declaration

Types in Souffle:  
number  
symbol (aka varchar)

Insert data

CSE 414 - Autumn 2018 9

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).
```

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

CSE 414 - Autumn 2018 10

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).
```

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

Find Movies made in 1940

CSE 414 - Autumn 2018 11

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).
```

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

SQL

```
SELECT name
FROM Movie
WHERE year = 1940
```

Find Movies made in 1940

CSE 414 - Autumn 2018 12

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

*id → name → year*

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

Order of variable matters!

Find Movies made in 1940

CSE 414 - Autumn 2018 13

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(idontCare,y,z),
z=1940.
```

Find Movies made in 1940

CSE 414 - Autumn 2018 14

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(_,y,z), z=1940.
```

\_ = "don't care" variables

Find Movies made in 1940

CSE 414 - Autumn 2018 15

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

CSE 414 - Autumn 2018 16

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

CSE 414 - Autumn 2018 17

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

CSE 414 - Autumn 2018 18

## Datalog: Terminology



f, l = head variables  
x, y, z = existential variables

CSE 414 - Autumn 2018

19

## More Datalog Terminology

$Q(\text{args}) \text{ :- } R_1(\text{args}), R_2(\text{args}), \dots$

- $R_i(\text{args}_i)$  called an *atom*, or a *relational predicate*
- $R_i(\text{args}_i)$  evaluates to true when relation  $R_i$  contains the tuple described by  $\text{args}_i$ .
  - Example: Actor(344759, 'Douglas', 'Fowley') is true
- In addition we can also have arithmetic predicates
  - Example:  $z > 1940$ .

Book uses AND instead of ,  $Q(\text{args}) \text{ :- } R_1(\text{args}) \text{ AND } R_2(\text{args}) \dots$

CSE 414 - Autumn 2018

20

## Datalog program


- A Datalog program consists of several rules
- Importantly, rules may be recursive!
  - Recall CSE 143!
- Usually there is one distinguished predicate that's the output
- We will show an example first, then give the general semantics.

CSE 414 - Autumn 2018

21

R encodes a graph  
e.g., connected cities

## Example



R=

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

CSE 414 - Autumn 2018


22

R encodes a graph  
e.g., connected cities

## Example

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

What does it compute?



R=

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

CSE 414 - Autumn 2018


23

R encodes a graph  
e.g., connected cities

## Example

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

What does it compute?



R=

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

Initially:  
T is empty.

CSE 414 - Autumn 2018

24

R encodes a graph e.g., connected cities

**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

First rule generates this

Second rule generates nothing (because T is empty)

CSE 414 - Autumn 2018 25

R encodes a graph e.g., connected cities

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

First rule generates this

Second rule generates this

CSE 414 - Autumn 2018 26

R encodes a graph e.g., connected cities

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

Both rules

First rule

Second rule

Third iteration: T =

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

New fact

CSE 414 - Autumn 2018 27

R encodes a graph e.g., connected cities

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

Fourth iteration: T = (same)

No new facts. DONE

CSE 414 - Autumn 2018 28

## Datalog Semantics

Fixpoint semantics

- Start:
  - $IDB_0$  = empty relations
  - $t = 0$
- Repeat:
  - $IDB_{t+1}$  = Compute Rules( $EDB, IDB_t$ )
  - $t = t+1$
- Until  $IDB_t = IDB_{t-1}$

CSE 414 - Autumn 2018 29

## More Features

- Aggregates
- Grouping
- Negation

CSE 414 - Autumn 2018 30

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

## Aggregates

`[aggregate name] <var> : { [relation to compute aggregate on] }`

`min x : { Actor(x, y, _) , y = 'John' }`

`Q(minId) :- minId = min x : { Actor(x, y, _) , y = 'John' }`

Assign variable to the value of the aggregate

Aggregates in Souffle:

- count
- min
- max
- sum

Meaning (in SQL)

`SELECT min(id) as minId  
FROM Actor as a  
WHERE a.name = 'John'`

CSE 414 - Autumn 2018 31

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

## Counting

`Q(c) :- c = count : { Actor(_, y, _) , y = 'John' }`

No variable here!

Meaning (in SQL, assuming no NULLs)

`SELECT count(*) as c  
FROM Actor as a  
WHERE a.name = 'John'`

CSE 414 - Autumn 2018 32

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

## Grouping

`Q(y,c) :- Movie(_,y), c = count : { Movie(_,y) }`

Meaning (in SQL)

`SELECT m.year, count(*)  
FROM Movie as m  
GROUP BY m.year`

CSE 414 - Autumn 2018 33

ParentChild(p,c)

## Examples

A genealogy database (parent/child)

```

graph TD
    Alice --> Carol
    Alice --> Bob
    Bob --> David
    Bob --> Eve
    Carol --> Eve
    Eve --> Fred
    Eve --> George
    Fred --> George
  
```

p	c
Alice	Carol
Bob	Carol
Bob	David
Carol	Eve
...	

CSE 414 - Autumn 2018 34

ParentChild(p,c)

## Count Descendants

For each person, count his/her descendants

CSE 414 - Autumn 2018 35

ParentChild(p,c)

## Count Descendants

For each person, count his/her descendants

Paths:

x	y
Alice	Carol
Alice	Eve
Alice	Fred
Alice	George

CSE 414 - Autumn 2018 36

ParentChild(p,c)

## Count Descendants

For each person, count his/her descendants

Paths:

x	y
Alice	Carol
Alice	Eve
Alice	Fred
Alice	George

Descendants:

x	count
Alice	4

CSE 414 - Autumn 2018 37

ParentChild(p,c)

## Count Descendants

For each person, count his/her descendants

Answer

p	cnt
Alice	4
Bob	5
Carol	3
David	2
Fred	1

CSE 414 - Autumn 2018 38

ParentChild(p,c)

## Count Descendants

For each person, count his/her descendants

Answer

p	cnt
Alice	4
Bob	5
Carol	3
David	2
Fred	1

Note: Eve and George do not appear in the answer (why?)

CSE 414 - Autumn 2018 41

ParentChild(p,c)

## Count Descendants

For each person, compute the total number of descendants

```
// for each person, compute his/her descendants
```

CSE 414 - Autumn 2018 40

ParentChild(p,c)

## Count Descendants

For each person, compute the total number of descendants

```
// for each person, compute his/her descendants
D(x,y) :- ParentChild(x,y).
```

CSE 414 - Autumn 2018 41

ParentChild(p,c)

## Count Descendants

For each person, compute the total number of descendants

```
// for each person, compute his/her descendants
D(x,y) :- ParentChild(x,y).
D(x,z) :- D(x,y), ParentChild(y,z).
```

CSE 414 - Autumn 2018 42

## Count Descendants

For each person, compute the total number of descendants

```
// for each person, compute his/her descendants
D(x,y) :- ParentChild(x,y).
D(x,z) :- D(x,y), ParentChild(y,z).

// For each person, count the number of descendants
```

## Count Descendants

For each person, compute the total number of descendants

```
// for each person, compute his/her descendants
D(x,y) :- ParentChild(x,y).
D(x,z) :- D(x,y), ParentChild(y,z).

// For each person, count the number of descendants
T(p,c) :- D(p,_), c = count : { D(p,y) }.
```

## Count Descendants

How many descendants does Alice have?

```
// for each person, compute his/her descendants
D(x,y) :- ParentChild(x,y).
D(x,z) :- D(x,y), ParentChild(y,z).

// For each person, count the number of descendants
T(p,c) :- D(p,_), c = count : { D(p,y) }.
```

## Count Descendants

How many descendants does Alice have?

```
// for each person, compute his/her descendants
D(x,y) :- ParentChild(x,y).
D(x,z) :- D(x,y), ParentChild(y,z).

// For each person, count the number of descendants
T(p,c) :- D(p,_), c = count : { D(p,y) }.

// Find the number of descendants of Alice
```

## Count Descendants

How many descendants does Alice have?


```
// for each person, compute his/her descendants
D(x,y) :- ParentChild(x,y).
D(x,z) :- D(x,y), ParentChild(y,z).

// For each person, count the number of descendants
T(p,c) :- D(p,_), c = count : { D(p,y) }.

// Find the number of descendants of Alice
Q(d) :- T(p,d), p = "Alice".
```

## Negation: use "!"


Find all descendants of Bob that are not descendants of Alice





### Negation: use “!”

Find all descendants of Bob that are not descendants of Alice



Answer

x
David

### Negation: use “!”

Find all descendants of Bob that are not descendants of Alice

```
// for each person, compute his/her descendants
D(x,y) :- ParentChild(x,y).
D(x,z) :- D(x,y), ParentChild(y,z).
```

### Negation: use “!”


Find all descendants of Bob that are not descendants of Alice

```
// for each person, compute his/her descendants
D(x,y) :- ParentChild(x,y).
D(x,z) :- D(x,y), ParentChild(y,z).

// Compute the answer: notice the negation
Q(x) :- D("Bob",x), !D("Alice",x).
```

### 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(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
:- 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)
```

## 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
```

ParentChild(p,c)

## Safe Datalog Rules

Here are unsafe datalog rules. What's "unsafe" about them ?

```
U1(x,y) :- ParentChild("Alice",x), y != "Bob"
```

```
U2(x) :- ParentChild("Alice",x), !ParentChild(x,y)
```

```
U3(minId, y) :- minId = min x : { Actor(x, y, _) }
```

CSE 414 - Autumn 2018 61

ParentChild(p,c)

## Safe Datalog Rules

Here are unsafe datalog rules. What's "unsafe" about them ?

```
U1(x,y) :- ParentChild("Alice",x), y != "Bob"
```

```
U2(x) :- ParentChild("Alice",x), !ParentChild(x,y)
```

```
U3(minId, y) :- minId = min x : { Actor(x, y, _) }
```

Hold's for every y other than "Bob"  
U1 = infinite!

CSE 414 - Autumn 2018 62

ParentChild(p,c)

## Safe Datalog Rules

Here are unsafe datalog rules. What's "unsafe" about them ?

```
U1(x,y) :- ParentChild("Alice",x), y != "Bob"
```

```
U2(x) :- ParentChild("Alice",x), !ParentChild(x,y)
```

```
U3(minId, y) :- minId = min x : { Actor(x, y, _) }
```

Hold's for every y other than "Bob"  
U1 = infinite!

Want Alice's childless children, but we get all children x (because there exists some y that x is not parent of y)

CSE 414 - Autumn 2018 63

ParentChild(p,c)

## Safe Datalog Rules

Here are unsafe datalog rules. What's "unsafe" about them ?

```
U1(x,y) :- ParentChild("Alice",x), y != "Bob"
```

```
U2(x) :- ParentChild("Alice",x), !ParentChild(x,y)
```

```
U3(minId, y) :- minId = min x : { Actor(x, y, _) }
```

Hold's for every y other than "Bob"  
U1 = infinite!

Want Alice's childless children, but we get all children x (because there exists some y that x is not parent of y)

Unclear what y is

CSE 414 - Autumn 2018 64

ParentChild(p,c)

## Safe Datalog Rules

Here are unsafe datalog rules. What's "unsafe" about them ?

```
U1(x,y) :- ParentChild("Alice",x), y != "Bob"
```

```
U2(x) :- ParentChild("Alice",x), !ParentChild(x,y)
```

```
U3(minId, y) :- minId = min x : { Actor(x, y, _) }
```

A datalog rule is safe if every variable appears in some positive, non-aggregated relational atom

CSE 414 - Autumn 2018 65

## Stratified Datalog

- Recursion does not cope well with aggregates or negation
- Example: what does this mean?
 

```
A() :- !B().
```

```
B() :- !A().
```
- A datalog program is stratified if it can be partitioned into *strata*
  - Only IDB predicates defined in strata 1, 2, ..., n may appear under ! or agg in stratum n+1.
- Many Datalog DBMSs (including souffle) accept only stratified Datalog.

CSE 414 - Autumn 2018 66

## Stratified Datalog

```

D(x,y) :- ParentChild(x,y).
D(x,z) :- D(x,y), ParentChild(y,z).
T(p,c) :- D(p,_), c = count : { D(p,y) }.
Q(d) :- T(p,d), p = "Alice".
    
```

Stratum 1

Stratum 2

May use D  
in an agg since it was  
defined in previous  
stratum

CSE 414 - Autumn 2018

67

## Stratified Datalog

```

D(x,y) :- ParentChild(x,y).
D(x,z) :- D(x,y), ParentChild(y,z).
T(p,c) :- D(p,_), c = count : { D(p,y) }.
Q(d) :- T(p,d), p = "Alice".
    
```

Stratum 1

Stratum 2

May use D  
in an agg since it was  
defined in previous  
stratum

```

D(x,y) :- ParentChild(x,y).
D(x,z) :- D(x,y), ParentChild(y,z).
Q(x) :- D("Alice",x), ID("Bob",x).
    
```

Stratum 1

Stratum 2

May use ID

```

A() :- !B().
B() :- !A().
    
```

Non-stratified

Cannot use !A

68

## Stratified Datalog

- If we don't use aggregates or negation, then the Datalog program is already stratified
- If we do use aggregates or negation, it is usually quite natural to write the program in a stratified way

CSE 414 - Autumn 2018

69