# CSE 544 <br> Principles of Database Management Systems <br> Lectures 6: Datalog (2) 

## Reminders

- This Friday: project proposals due (turnin using git)
- Monday: paper review due (12h before lecture)
- Next Friday: brief meetings to discuss your project
- Next Friday: hw2 due


## Suggested Readings for Datalog

- Joe Hellerstein, "The Declarative Imperative," SIGMOD Record 2010
- R\&G Chapter 24
- Phokion Kolaitis' tutorial on database theory at Simon's https://simons.berkeley.edu/sites/default/files/docs/5241/ simons16-21.pdf
- Daniel Zinn, Todd J. Green, Bertram Ludäscher: Winmove is coordination-free (sometimes). ICDT 2012


## Review

- What is datalog?
- What is the naïve evaluation algorithm?
- What is the seminaive algorithm?


## Outline

## - Semi-joins

- Semi-join reduction
- Acyclic queries
- Magic sets


## Cost of Computing a Query

- Suppose $|\mathrm{R}|=|\mathrm{S}|=\mathrm{n}$
- What is the cost of a join $R \bowtie S$ ?

$$
q(x, y, z)=R(x, y), S(y, z)
$$

- Algorithms (discuss in class):


## Cost of Computing a Query

- Suppose $|\mathrm{R}|=|\mathrm{S}|=\mathrm{n}$
- What is the cost of a join $R \bowtie S$ ?

$$
q(x, y, z)=R(x, y), S(y, z)
$$

- Algorithms (discuss in class):
- Nested loop join
- Hash join
- Merge join


## Cost of Computing a Query

- Suppose $|\mathrm{R}|=|\mathrm{S}|=\mathrm{n}$
- What is the cost of a join $R \bowtie S$ ?

$$
q(x, y, z)=R(x, y), S(y, z)
$$

- Algorithms (discuss in class):
- Nested loop join
- Hash join
- Merge join

$O(n) \ldots O\left(n^{2}\right)$
$\mathrm{O}(\mathrm{n} \log \mathrm{n}) \ldots \mathrm{O}\left(\mathrm{n}^{2}\right)$



## Cost of Computing a Query

- Suppose $|\mathrm{R}|=|\mathrm{S}|=|\mathrm{T}|=|\mathrm{K}|=\mathrm{n}$
- What is the complexity of computing these queries?

$$
\begin{equation*}
\text { Q1 }(x, y, z)=R(x, y), S(y, z) \tag{2}
\end{equation*}
$$

## Cost of Computing a Query

- Suppose $|\mathrm{R}|=|\mathrm{S}|=|\mathrm{T}|=|\mathrm{K}|=\mathrm{n}$
- What is the complexity of computing these queries?

$$
\begin{equation*}
\text { Q1 }(x, y, z)=R(x, y), S(y, z) \tag{2}
\end{equation*}
$$

Q2 $(x, y, z, u)=R(x, y), S(y, z), T(z, u)$

## Cost of Computing a Query

- Suppose $|\mathrm{R}|=|\mathrm{S}|=|\mathrm{T}|=|\mathrm{K}|=\mathrm{n}$
- What is the complexity of computing these queries?

$$
\begin{align*}
& \mathrm{Q} 1(\mathrm{x}, \mathrm{y}, \mathrm{z})=\mathrm{R}(\mathrm{x}, \mathrm{y}), \mathrm{S}(\mathrm{y}, \mathrm{z})  \tag{2}\\
& \mathrm{Q} 2(\mathrm{x}, \mathrm{y}, \mathrm{z}, \mathrm{u})=\mathrm{R}(\mathrm{x}, \mathrm{y}), \mathrm{S}(\mathrm{y}, \mathrm{z}), \mathrm{T}(\mathrm{z}, \mathrm{u}) \tag{3}
\end{align*}
$$

## Cost of Computing a Query

- Suppose $|\mathrm{R}|=|\mathrm{S}|=|\mathrm{T}|=|\mathrm{K}|=\mathrm{n}$
- What is the complexity of computing these queries?

$$
\begin{align*}
& \text { Q1 }(x, y, z)=R(x, y), S(y, z)  \tag{2}\\
& \text { Q2(x,y,z,u) }=R(x, y), S(y, z), T(z, u)  \tag{3}\\
& \text { Q3(x,y,z,u,v) }=R(x, y), S(y, z), T(z, u), K(u, v)
\end{align*}
$$

## Cost of Computing a Query

- Suppose $|\mathrm{R}|=|\mathrm{S}|=|\mathrm{T}|=|\mathrm{K}|=\mathrm{n}$
- What is the complexity of computing these queries?

$$
\begin{align*}
& \text { Q1 }(x, y, z)=R(x, y), S(y, z)  \tag{2}\\
& \text { Q2(x,y,z,u) }=R(x, y), S(y, z), T(z, u)  \tag{3}\\
& \text { Q3(x,y,z,u,v) }=R(x, y), S(y, z), T(z, u), K(u, v) \tag{4}
\end{align*}
$$

## Cost of Computing a Query

- Suppose $|\mathrm{R}|=|\mathrm{S}|=|\mathrm{T}|=|\mathrm{K}|=\mathrm{n}$
- What is the complexity of computing these queries?

$$
\begin{align*}
& \text { Q1 }(x, y, z)=R(x, y), S(y, z)  \tag{2}\\
& \text { Q2(x,y,z,u) }=R(x, y), S(y, z), T(z, u)  \tag{3}\\
& \text { Q3(x,y,z,u,v) }=R(x, y), S(y, z), T(z, u), K(u, v) \tag{4}
\end{align*}
$$

Ideally cost: O(|lnput| + |Output|)

## Cost of Computing a Query

- Naïve computation often exceeds this bound
- $Q(x, y, z, u)=R\left({ }^{\prime} a^{\prime}, y\right), S(y, z), T(z, u), K\left(u,{ }^{\prime}{ }^{\prime}\right)$



## Cost of Computing a Query

- Naïve computation often exceeds this bound
- $Q(x, y, z, u)=R\left({ }^{(a ', ~ y), ~ S(y, z), T(z, u), K\left(u,{ }^{\prime}{ }^{\prime}\right) ~}\right.$

$$
\begin{aligned}
& R=\left\{a^{\prime}\right\} \times\{1, \ldots, n / 2\} \\
& S=\{1, \ldots, n / 2\} \times\left\{a^{\prime}\right\} \cup\{n / 2+1, \ldots, n\} \times\left\{b^{\prime}\right\} \\
& T=\left\{a^{\prime}\right\} \times\{1, \ldots, n / 2\} \cup\left\{b^{\prime}\right\} \times\{n / 2+1, \ldots, n\} \\
& K=\{n / 2+1, \ldots, n\} \times\left\{b^{\prime}\right\}
\end{aligned}
$$



## Cost of Computing a Query

- Naïve computation often exceeds this bound
- $Q(x, y, z, u)=R\left({ }^{\prime} a^{\prime}, y\right), S(y, z), T(z, u), K\left(u,{ }^{\prime} b^{\prime}\right)$

$$
\begin{aligned}
& R=\left\{a^{\prime}\right\} \times\{1, \ldots, n / 2\} \\
& S=\{1, \ldots, n / 2\} \times\left\{a^{\prime}\right\} \cup\{n / 2+1, \ldots, n\} \times\left\{b^{\prime}\right\} \\
& T=\left\{a^{\prime}\right\} \times\{1, \ldots, n / 2\} \cup\left\{b^{\prime}\right\} \times\{n / 2+1, \ldots, n\} \\
& K=\{n / 2+1, \ldots, n\} \times\left\{b^{\prime}\right\}
\end{aligned}
$$




## The Semijoin Operator

## Definition: the semi-join operation is $R \ltimes S=\Pi_{\text {Attr R })}(R \bowtie S)$

## Properties of Semijoins

- $R(A, B) \ltimes S(B, C)$ same as $Q(A, B):-R(A, B), S(B, C)$
- Cost: $\mathrm{O}(|\mathrm{R}|+|\mathrm{S}|)$ (ignoring log-factors)
- Cost is independent on the join output
- The law of semijoins is:

$$
R \bowtie S=(R \ltimes S) \bowtie S
$$

Consequence: we can perform a semi-join before a join

## Outine

- Semi-joins
- Semi-join reduction
- Acyclic queries
- Magic sets


## Semijoin Optimizations

- In parallel databases: often combined with Bloom Filters (pp. 747 in the textbook)
- Magic sets for datalog were invented after semi-join reductions, and the connection became clear only later
- Some complex semi-join reductions for non-recursive SQL optimizations are sometimes called "magic sets"


## Semijoin Reducer

- Given a query:

$$
Q=R_{1} \bowtie R_{2} \bowtie \ldots \bowtie R_{n}
$$

- A semijoin reducer for Q is

$$
\begin{aligned}
& R_{i 1}=R_{i 1} \ltimes R_{j 1} \\
& R_{i 2}=R_{i 2} \ltimes R_{j 2} \\
& \cdots \cdots=R_{i p} \ltimes R_{j p} \\
& \mathrm{R}_{\mathrm{ip}}={ }_{2}
\end{aligned}
$$

such that the query is equivalent to:

$$
Q=R_{k 1} \bowtie R_{k 2} \bowtie \ldots \bowtie R_{k n}
$$

- A full reducer is such that no dangling tuples remain


## Example

- Example:

$$
Q=R(A, B) \bowtie S(B, C)
$$

- A semijoin reducer is:

$$
R_{1}(A, B)=R(A, B) \ltimes S(B, C)
$$

- The rewritten query is:

$$
Q=R_{1}(A, B) \bowtie S(B, C)
$$

## Semijoin Reducer

- More complex example:

$$
Q(y, z, u)=R\left({ }^{\prime} a^{\prime}, y\right), S(y, z), T(z, u), K\left(u, b^{\prime}\right)
$$

- Find a full reducer


## Semijoin Reducer

- More complex example:

$$
Q(y, z, u)=R\left(a^{\prime}, y\right), S(y, z), T(z, u), K\left(u, b^{\prime}\right)
$$

- Find a full reducer


## Semijoin Reducer

- More complex example:

$$
Q(y, z, u)=R\left(a^{\prime}, y\right), S(y, z), T(z, u), K\left(u, b^{\prime}\right)
$$

- Find a full reducer

$$
\begin{aligned}
& S^{\prime}(y, z):-S(y, z) \ltimes R\left(a^{\prime}, y\right) \\
& T^{\prime}(z, u):-T(z, u) \ltimes S^{\prime}(y, z) \\
& K^{\prime}(u):-K\left(u,{ }^{\prime} b^{\prime}\right) \ltimes T^{\prime}(z, u) \\
& T^{\prime \prime}(z, u):-T^{\prime}(z, u) \ltimes K^{\prime}(u) \\
& S^{\prime \prime}(y, z):-S^{\prime}(y, z) \ltimes T^{\prime \prime}(z, u) \\
& \text { R" }(y):-R\left({ }^{\prime} a^{\prime}, y\right) \ltimes S^{\prime \prime}(y, z)
\end{aligned}
$$

## Semijoin Reducer

- More complex example:

$$
Q(y, z, u)=R\left({ }^{\prime} a^{\prime}, y\right), S(y, z), T(z, u), K\left(u, b^{\prime}\right)
$$

- Find a full reducer

$$
\begin{aligned}
& S^{\prime}(y, z):-S(y, z) \ltimes R\left({ }^{\prime} a^{\prime}, y\right) \\
& T^{\prime}(z, u):-T(z, u) \ltimes S^{\prime}(y, z) \\
& K^{\prime}(u):-K\left(u,{ }^{\prime} b^{\prime}\right) \ltimes T^{\prime}(z, u) \\
& T^{\prime \prime}(z, u):-T^{\prime}(z, u) \ltimes K^{\prime}(u) \\
& S^{\prime \prime}(y, z):-S^{\prime}(y, z) \ltimes T^{\prime \prime}(z, u) \\
& R^{\prime \prime}(y):-R\left({ }^{\prime} a^{\prime}, y\right) \ltimes S^{\prime \prime}(y, z)
\end{aligned}
$$

- Finally, compute:

$$
Q(y, z, u)=R^{\prime \prime}(y), S^{\prime \prime}(y, z), T^{\prime \prime}(z, u), K^{\prime \prime}(u)
$$

## Practice at Home...

- Find semi-join reducer for $R(x, y), S(y, z), T(z, u), K(u, v), L(v, w)$


## Not All Queries Have Full Reducers

- Example:

$$
Q=R(A, B) \bowtie S(B, C) \bowtie T(A, C)
$$

- Can write many different semi-join reducers
- But no full reducer of length $\mathrm{O}(1)$ exists


## Outline

- Semi-joins
- Semi-join reduction
- Acyclic queries
- Magic sets


## Acyclic Queries

- Fix a Conjunctive Query without self-joins
- $Q$ is acyclic if its atoms can be organized in a tree such that for every variable the set of nodes that contain that variable form a connected component

$R(x, y), S(y, z), T(z, x)$ is cyclic


## Yannakakis Algorithm

- Given: acyclic query Q
- Compute Q on any database in time $O(|\ln n u t|+|O u t p u t|)$
- Step 1: semi-join reduction
- Pick any root node $x$ in the tree decomposition of $Q$
- Do a semi-join reduction sweep from the leaves to $x$
- Do a semi-join reduction sweep from $x$ to the leaves
- Step 2: compute the joins bottom up, with early projections


## Examples in Class

- Boolean query: $Q()$ :- ...
- Non-boolean: $\mathrm{Q}(\mathrm{x}, \mathrm{m})$ :- ...
- With aggregate: $\mathrm{Q}(\mathrm{x}, \mathrm{sum}(\mathrm{m}))$ :- ...
- And also: $\mathrm{Q}\left(\mathrm{x}, \operatorname{count}\left({ }^{*}\right)\right)$ :- ...


In all cases: runtime $=\mathrm{O}(|\mathrm{R}|+|\mathrm{S}|+\ldots+\mid$ ㄴ $|+|$ Output $\mid)$

## Testing if $Q$ is Acyclic

- An ear of $Q$ is an atom $R(X)$ with the following property:
- Let $X^{\prime} \subseteq X$ be the set of join variables (meaning: they occur in at least one other atom)
- There exists some other atom $S(Y)$ such that $X^{\prime} \subseteq Y$
- The GYO algorithm (Graham,Yu,Özsoyoğlu) for testing if $Q$ is acyclic:
- While $Q$ has an ear $R(X)$, remove the atom $R(X)$ from the query
- If all atoms were removed, then $Q$ is acyclic
- If atoms remain but there is no ear, then $Q$ is cyclic
- Show example in class


## Outine

- Semi-joins
- Semi-join reduction
- Acyclic queries
- Magic sets


## Magic Sets

- Problem: datalog programs compute a lot, but sometimes we need only very little
- Prolog computes top-down and retrieves very little datalog computes bottom up retrieves a lot
- (Prolog has other issues: left recursive prolog never terminates!)
- Magic sets transform a datalog program P into a new program $\mathrm{P}^{\prime}$, such that bottom-up( $\left.\mathrm{P}^{\prime}\right)=$ top-down $(\mathrm{P})$


## Example 1



Bottom-up evaluation very inefficient

## Example 1



$$
\begin{aligned}
& \mathrm{T}(\mathrm{x}, \mathrm{y}):-\mathrm{E}(\mathrm{x}, \mathrm{y}) \\
& \mathrm{T}(\mathrm{x}, \mathrm{y}):-\mathrm{T}(\mathrm{x}, \mathrm{z}), \mathrm{E}(\mathrm{z}, \mathrm{y}) \\
& \mathrm{Q}(\mathrm{y}):-\mathrm{T}(3, y) \\
& \hline
\end{aligned}
$$

Manual optimization:

$$
\begin{aligned}
& Q(y):-E(3, y) \\
& Q(y):-Q(x), E(x, y)
\end{aligned}
$$

## Example 2



## Same generation

## SG(x,x) :- V(x) SG(x,y) :- Up(x,u),SG(u,v),Dn(u,y) $Q(y):-S G(1, y)$

Manual optimization???


## Magic Set Rewriting (simplified)

- For each IDB predicate create "adorned" versions, with binding patters
- For each adorned IDB P, create a predicate Magic ${ }_{P}$
- For each rule, create several rules, one for each possible adornment of the head:
- Allow information to flow left-to-right ("sideways information passing"), and this defines the required adornments of the IDB's in the body
- If there are $k$ IDB's in the body, create $k+1$ supplementary relations Supp ${ }_{i}$, which guard the set of bound variables passed on to the i'th IDB
- New rules defining Magic ${ }_{p}$ : one for the query, and one for each Supp ${ }_{i}$ preceding an occurrence of $P$ in a body


## Adorned predicate

- $b=b o u n d, f=f r e e$
- $T^{b f}(x, y)$ means:
- The values of $x$ are known
- The values of $y$ are not known (need to be retrieved)
- Need to create all combinations: $\mathrm{T}^{\mathrm{bf}}, \mathrm{T}^{\mathrm{fb}}$
- Side-ways information passing means that we adorn rules allowing information to flow left-to-right
- E.g. $\quad T(x, y):-E(x, u), T(u, v), E(v, w), T(w, z), E(z, y)$
- Adorned: $T^{b f}(x, y):-E(x, u), T^{b f}(u, v), E(v, w), T^{b f}(w, z), E(z, y)$


## Supplementary Relations

- Given adornment $T^{b f}(x, y)$, a new predicate $\operatorname{Supp}(x)$ contains the (small!) set of values $x$ for which we want to compute $\mathrm{T}^{\mathrm{bf}}(\mathrm{x}, \mathrm{y})$
- E.g. $T^{b f}(x, y):-E(x, u), T^{b f}(u, v), E(v, w), T^{b f}(w, z), E(z, y)$



## Supp Rules

- E.g. $T^{b f}(x, y):-E(x, u), T^{b f}(u, v), E(v, w), T^{b f}(w, z), E(z, y)$


Becomes:

- $\operatorname{Supp}_{0}(\mathrm{x}):-\operatorname{Magic}_{\mathrm{Tbf}}(\mathrm{x})$ /* next slide ... */
- $\operatorname{Supp}_{1}(\mathrm{x}, \mathrm{u}):-\operatorname{Supp}_{0}(\mathrm{x}), \mathrm{E}(\mathrm{x}, \mathrm{u})$
- $\operatorname{Supp}_{2}(\mathrm{x}, \mathrm{w}):-\operatorname{Supp}_{1}(\mathrm{x}, \mathrm{u}), \mathrm{T}^{\mathrm{bf}}(\mathrm{u}, \mathrm{v}), \mathrm{E}(\mathrm{v}, \mathrm{w})$
- $\operatorname{Supp}_{3}(x, y):-\operatorname{Supp}_{2}(x, w), T^{b f}(w, z), E(z, y)$
- $\mathrm{Tb}^{\mathrm{bf}}(\mathrm{x}, \mathrm{y}):-\operatorname{Supp}_{3}(\mathrm{x}, \mathrm{y})$


## Adding the Magic Predicate

- E.g. $T^{b f}(x, y):-E(x, u), T^{b f}(u, v), E(v, w), T^{b f}(w, z), E(z, y)$

- Magic $_{\text {Tbf }}(x)=$ the set of bounded values of $x$ for which we need to compute $\mathrm{T}^{\mathrm{bf}}(\mathrm{x}, \mathrm{y})$
- E.g.
- Magic $_{\text {Tbf }}(3)$ :- $\quad / *$ if the query is $Q(y):-T(3, y)$ */
- Magic $_{\text {Tbf }}(\mathrm{u}):-\operatorname{Supp}_{1}(\mathrm{x}, \mathrm{u}) /^{*}$ need to compute $\mathrm{T}^{\mathrm{bf}}(\mathrm{u}, \mathrm{v}) * /$
- $\operatorname{Magic}_{\text {Tbf }}(\mathrm{w}):-\operatorname{Supp}_{2}(\mathrm{x}, \mathrm{w}) /^{*}$ need to compute $\mathrm{T}^{\mathrm{bf}}(\mathrm{w}, \mathrm{z})^{* /}$


## Example 1



## Magic Sets

Original:

$$
\begin{aligned}
& T(x, y):-E(x, y) \\
& T(x, y):-T(x, z), E(z, y) \\
& Q(y):-T(3, y)
\end{aligned}
$$

Adorned:

## Example 1



Original:

$$
\begin{aligned}
& T(x, y):-E(x, y) \\
& T(x, y):-T(x, z), E(z, y) \\
& Q(y):-T(3, y)
\end{aligned}
$$

Adorned:

```
```

Tbf}(x,y):- E(x,y

```
```

Tbf}(x,y):- E(x,y
Tbf(x,y):- Tbf(x,z),E(z,y)
Tbf(x,y):- Tbf(x,z),E(z,y)
Q(y) :- Tbf

```
```

Q(y) :- Tbf

```
```


## Magic Sets

## Example 1



Original:

$$
\begin{aligned}
& T(x, y):-E(x, y) \\
& T(x, y):-T(x, z), E(z, y) \\
& Q(y):-T(3, y)
\end{aligned}
$$

Adorned:

$$
\begin{aligned}
& \mathrm{T}^{\mathrm{bf}}(\mathrm{x}, \mathrm{y}):-\mathrm{E}(\mathrm{x}, \mathrm{y}) \\
& \mathrm{T}^{\mathrm{bf}}(\mathrm{x}, \mathrm{y}):-\mathrm{T}^{\mathrm{bf}}(\mathrm{x}, \mathrm{z}), \mathrm{E}(\mathrm{z}, \mathrm{y}) \\
& \mathrm{Q}(\mathrm{y}):-\mathrm{T}^{\mathrm{bf}}(3, \mathrm{y})
\end{aligned}
$$

## Magic Sets

```
/* T(x,y) :- E(x,y) */
Supp
Supp
Tbf(x,y):- Supp
/* T(x,y) :- T(x,z),E(z,y) */
Supp'0(x) :- - Magic
Supp'(x,z) :- Supp'0}(x), Tbf(x,z
Supp'(x,y):- Supp',}(x,z), E(z,y
Tbf}(x,y):- Supp'2(x,y
/* Q(y) :- T(3,y) */
Magic
Magic
```


## Practice at home



$$
\begin{aligned}
& T(x, y):-E(x, y) \\
& T(x, y):-T(x, z), T(z, y) \\
& Q(y):-T(3, y)
\end{aligned}
$$

