# Query Optimization 

 Exercise Session 4Andrey Gubichev

November 10, 2014

## Homework: Selectivity estimations

The selectivity of $\sigma_{R 1 . x=c}$ is...

- if $x$ is the key: $\frac{1}{|R 1|}$
- if $x$ is not the key: $\frac{1}{|R 1 . x|}$

The selectivity of $\bowtie_{R 1 . x=R 2 . y}$ is...

- if both $x$ and $y$ are the keys: $\frac{1}{\max (|R 1|,|R 2|)}$
- if only $x$ is the key: $\frac{1}{|R 1|}$
- if both $x$ and $y$ are not the keys: $\frac{1}{\max (|R 1 . x|,|R 2 . y|)}$

When our selectivity estimations are bad?

## Example 1.1, due Bernhard Radke

```
select *
from title t, movie_companies mc,
    company_name n
where mc.movie_id = t.id
    and mc.company_id=n.id
    and n.country_code='[us]'
    and t.production_year=1880
```


## Example 1.1, due Atanas Mirchev

select m2.movie_id
from movie_info m1, movie_link l, movie_info m2
where m1.movie_id = l.movie_id
and m2.movie_id = l.linked_movie_id
m2.info = 'Comedy' and m1.info = 'Lifestyle'

## Example 1.2, due André Kohn, Valentin Hauner

```
select * from title t,
    movie_info mi, kind_type kt
where t.id = mi.movie_id
    and t.kind_id = kt.id
    and t.title like '%from%'
```


## Example 1.2, due Michele Bertoni

Movies with certain names, actors with certain names that acted as 'guests'

SELECT *
FROM role_type r
join cast_info c1 on r.id = c1.role_id
join name n1 on c1.person_id = n1.id
join cast_info c2 on c2.movie_id = c1.movie_id
join name n2 on n2.id $=$ c2.person_id
WHERE n1.name like , $\% \mathrm{~J} \% \mathrm{~J} \%$,
and n2.name like '\%J\%J\%,
and r.role = 'guest';

## Operator Optimality

When is the nested loops join optimal?

## Operator Optimality

When is the nested loops join optimal?

```
select *
```

from $R$ join $S$ on R.str=S.str
where R.ID = 1

- $R$ has the key $I D$


## Operator Optimality

When is the blocked nested loops join optimal?

## Operator Optimality

When is the blocked nested loops join optimal?

```
select R.b, S.b
from R, S
where R.b <> S.b * S.c
```


## Operator Optimality

When is the merge join optimal?

## Operator Optimality

When is the merge join optimal?

```
select *
from R, S
where R.a = S.b
order by R.a asc
```


## Greedy operator ordering

- take the query graph
- find relations $R_{1}, R_{2}$ such that $\left|R_{1} \bowtie R_{2}\right|$ is minimal and merge them into one node
- repeat until the query graph has more than one node

Generates bushy trees!

## Example



## Example - step 1



## Example - after step 1



## Example - step 2



## Example- after step 2



## Example - step 3



## Example - after step 3



## Example - step 4

$$
\begin{gathered}
R_{1}(10) \xrightarrow{0.8} R_{2}(10) \underline{0.5} R_{3} \bowtie R_{4}(30) \\
\left(R_{5} \bowtie R_{6}\right) \bowtie R_{9}(60) \\
R_{7} \bowtie R_{8}(30)
\end{gathered}
$$

## Example - after step 4



## Example - step 5



## Example - after step 5

$\left(R_{1} \bowtie R_{2}\right) \bowtie\left(R_{3} \bowtie R_{4}\right)(1200)$
0.2268
$\left(R_{7} \bowtie R_{8}\right) \bowtie\left(\left(R_{5} \bowtie R_{6}\right) \bowtie R_{9}\right)(1080)$

## Example - result



## IKKBZ (informally)

Query graph $Q$ is acyclic. Pick a root node, turn it into a tree. Run the following procedure for every root node, select the cheapest plan

Input: rooted tree $Q$

1. if the tree is a single chain, stop
2. find the subtree (rooted at $r$ ) all of whose children are chains
3. normalize, if $c_{1} \rightarrow c_{2}$, but $\operatorname{rank}\left(c_{1}\right)>\operatorname{rank}\left(c_{2}\right)$ in the subtree rooted at $r$
4. merge chains in the subtree rooted at $r$, rank is ascending
5. repeat 1

## IKKBZ (informally)

For every relation $R_{i}$ we keep

- cardinality $n_{i}$
- selectivity $s_{i}$ — the selectivity of the incoming edge from the parent of $R_{i}$
- cost $C\left(R_{i}\right)=n_{i} s_{i}$ (or 0 , if $R_{i}$ is the root)
- rank $r_{i}=\frac{n_{i} s_{i}-1}{n_{i} s_{i}}$

Moreover,

- $C\left(S_{1} S_{2}\right)=C\left(S_{1}\right)+T\left(S_{1}\right) C\left(S_{2}\right)$
- $T(S)=\prod_{R_{i} \in S}\left(s_{i} n_{i}\right)$
- rank of a sequence $r(S)=\frac{T(S)-1}{C(S)}$


## Understanding IKKBZ

- what is the rank?
- when is $\left(R_{1} \bowtie R_{2}\right) \bowtie R_{3}$ cheaper than $\left(R_{1} \bowtie R_{3}\right) \bowtie R_{2}$ ?


## Understanding IKKBZ

- what is the rank?
- when is $\left(R_{1} \bowtie R_{2}\right) \bowtie R_{3}$ cheaper than $\left(R_{1} \bowtie R_{3}\right) \bowtie R_{2}$ ?
- if $r\left(R_{2}\right)<r\left(R_{3}\right)$ !


## IKKBZ - example



## IKKBZ - example

Subtree $R_{3}$ : merging, $\operatorname{rank}\left(R_{5}\right)<\operatorname{rank}\left(R_{4}\right)$


| Relation | n | s | C | T | rank |
| :--- | :---: | :---: | :---: | :---: | :---: |
| 2 | 20 | $\frac{1}{5}$ | 4 | 4 | $\frac{3}{4}$ |
| 3 | 30 | $\frac{1}{3}$ | 10 | 10 | $\frac{9}{10}$ |
| 4 | 50 | $\frac{1}{10}$ | 5 | 5 | $\frac{4}{5}$ |
| 5 | 2 | 1 | 2 | 2 | $\frac{1}{2}$ |

## IKKBZ - example

Subtree $R_{1}$ :
$\operatorname{rank}\left(R_{3}\right)>\operatorname{rank}\left(R_{5}\right)$, normalizing


| Relation | n | s | C | T | rank |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 2 | 20 | $\frac{1}{5}$ | 4 | 4 | $\frac{3}{4}$ |
| 3 | 30 | $\frac{1}{3}$ | 10 | 10 | $\frac{9}{10}$ |
| 4 | 50 | $\frac{1}{10}$ | 5 | 5 | $\frac{4}{5}$ |
| 5 | 2 | 1 | 2 | 2 | $\frac{1}{2}$ |
| 3,5 | 60 | $\frac{1}{3}$ | 30 | 20 | $\frac{19}{30}$ |

## IKKBZ - example

Subtree $R_{1}$ : merging


| Relation | n | s | C | T | rank |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 2 | 20 | $\frac{1}{5}$ | 4 | 4 | $\frac{3}{4}$ |
| 3 | 30 | $\frac{1}{15}$ | 10 | 10 | $\frac{9}{10}$ |
| 4 | 50 | $\frac{1}{10}$ | 5 | 5 | $\frac{4}{5}$ |
| 5 | 2 | 1 | 2 | 2 | $\frac{1}{2}$ |
| 3,5 | 60 | $\frac{1}{15}$ | 30 | 20 | $\frac{19}{30}$ |

## IKKBZ - example

Denormalizing
$R_{1}$


| Relation | n | s | C | T | rank |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 2 | 20 | $\frac{1}{5}$ | 4 | 4 | $\frac{3}{4}$ |
| 3 | 30 | $\frac{1}{15}$ | 10 | 10 | $\frac{9}{10}$ |
| 4 | 50 | $\frac{1}{10}$ | 5 | 5 | $\frac{4}{5}$ |
| 5 | 2 | 1 | 2 | 2 | $\frac{1}{2}$ |
| 3,5 | 60 | $\frac{1}{3}$ | 30 | 20 | $\frac{3}{30}$ |

## IKKBZ - another example



## IKKBZ



- $r\left(R_{2}\right)=\frac{9}{10}$
- $r\left(R_{3}\right)=\frac{4}{5}$
- $r\left(R_{4}\right)=0$
- $r\left(R_{5}\right)=\frac{13}{15}$
- $r\left(R_{6}\right)=\frac{9}{10}$
- $r\left(R_{7}\right)=\frac{4}{5}$
- $r\left(R_{8}\right)=\frac{19}{20}$
- $r\left(R_{9}\right)=\frac{3}{4}$


## IKKBZ



## IKKBZ



## IKKBZ

$$
R_{R_{2}}^{R_{1}-\frac{1}{10}} R_{3} \stackrel{\frac{1}{20}}{\frac{1}{10}} R_{4} \stackrel{\frac{3}{4}}{4} R_{5} \stackrel{\frac{1}{2}}{2} R_{8,9} \frac{\frac{1}{5}}{5} R_{6,7} \vee r\left(R_{8,9}\right)<r\left(R_{6,7}\right)
$$

## IKKBZ

> - $r\left(R_{2}\right)=\frac{9}{10}$
> - $r\left(R_{3}\right)=\frac{4}{5}$
> - $r\left(R_{4}\right)=0$

## IKKBZ



## IKKBZ



- $n_{5,8,9}=800$
- $C_{5,8,9}=\frac{1515}{2}$
- $T_{5,8,9}=600$
- $r\left(R_{5,8,9}\right)=\frac{1198}{1515}=$ 0.79..
- $r\left(R_{6,7}\right)=0.816 .$.


## IKKBZ

$$
\begin{aligned}
& \\
& r\left(R_{2}\right)=\frac{9}{10} \\
& \bullet r\left(R_{5,8,9}\right)=\frac{1198}{1515}= \\
& R_{1} \frac{1}{9} \\
& 0.79 . . \\
& R_{2} \frac{1}{10} r\left(R_{3}\right)=0.8 \\
& \\
& r\left(R_{4}\right)=0 \\
& \bullet r\left(R_{6,7}\right)=0.816 . .
\end{aligned}
$$

IKKBZ
$R_{1}-R_{3}-R_{4}-R_{5,8,9}-R_{6,7}-R_{2}$

IKKBZ
$R_{1}-R_{3}-R_{4}-R_{5}-R_{8}-R_{9}-R_{6}-R_{7}-R_{2}$

## IKKBZ-based heuristics

What if $Q$ has cycles?

- Observation 1: the answer of the query, corresponding to any subgraph of the query graph, is a superset of the answer to the original query
- Observation 2: a very selective join is more likely to be influential in choosing the order than a non-selective join


## IKKBZ-based heuristics

What if $Q$ has cycles?

- Observation 1: the answer of the query, corresponding to any subgraph of the query graph, is a superset of the answer to the original query
- Observation 2: a very selective join is more likely to be influential in choosing the order than a non-selective join

Choose the minimum spanning tree (minimize the product of the edge weights), compute the total order, compute the original query.

## Homework: Task 1 (15 points)

- Give an example query qraph with join selectivities for which the greedy operator ordering (GOO) algorithm does not give the optimal (with regards to $C_{\text {out }}$ ) join tree. Specify the optimal join tree. No cross products
- For that example perform the IKKBZ-based heuristics


## Homework: Task 2 (15 points)

- Using the program from the the last exercise as basis, construct the query graph for each connected component.


## Info

- Exercises due: 9 AM, November 17, 2014

