Query Optimization

Exercise Session 3

Andrey Gubichev

November 3, 2014

Homework: Task 1

```
select *
 from lineitem l, orders o, customers c
 where l.l_orderkey=o.o_orderkey
           and o.o_custkey=c.c_custkey
           and c.c_name='Customer#000014993'.
```

Homework: Task 2

We know $|R1|$, $|R2|$, domains of $R1.x$, $R2.y$,

Homework: Task 2

We know $|R1|$, $|R2|$, domains of $R1.x$, $R2.y$, (that is, $|R1.x|$, $|R2.y|$), and whether x and y are keys or not.
The selectivity of $\sigma_{R1.x=c}$ can be estimated as...

- if x is the key:

Homework: Task 2

We know $|R1|$, $|R2|$, domains of $R1.x$, $R2.y$, (that is, $|R1.x|$, $|R2.y|$), and whether x and y are keys or not.
The selectivity of $\sigma_{R1.x=c}$ can be estimated as...

- if x is the key: $\frac{1}{|R1|}$

Homework: Task 2

We know $|R1|$, $|R2|$, domains of $R1.x$, $R2.y$, (that is, $|R1.x|$, $|R2.y|$), and whether x and y are keys or not.
The selectivity of $\sigma_{R1.x=c}$ can be estimated as...

- if x is the key: $\frac{1}{|R1|}$
- if x is not the key:

Homework: Task 2

We know $|R1|$, $|R2|$, domains of $R1.x$, $R2.y$, (that is, $|R1.x|$, $|R2.y|$), and whether x and y are keys or not.
The selectivity of $\sigma_{R1.x=c}$ can be estimated as...

- if x is the key: $\frac{1}{|R1|}$
- if x is not the key: $\frac{1}{|R1.x|}$

Homework: Task 2

We know $|R1|$, $|R2|$, $|R1.x|$, $|R2.y|$, and whether x and y are keys or not.
First, the size of $R1 \times R2$ is

Homework: Task 2

We know $|R1|$, $|R2|$, $|R1.x|$, $|R2.y|$, and whether x and y are keys or not.
First, the size of $R1 \times R2$ is $|R1||R2|$
The selectivity of $\bowtie_{R1.x=R2.y}$ can be estimated as...

- if both x and y are the keys:

Homework: Task 2

We know $|R1|$, $|R2|$, $|R1.x|$, $|R2.y|$, and whether x and y are keys or not.
First, the size of $R1 \times R2$ is $|R1||R2|$
The selectivity of $\bowtie_{R1.x=R2.y}$ can be estimated as...

- if both x and y are the keys: $\frac{1}{max(|R1|,|R2|)}$

Homework: Task 2

We know $|R1|$, $|R2|$, $|R1.x|$, $|R2.y|$, and whether x and y are keys or not.
First, the size of $R1 \times R2$ is $|R1||R2|$
The selectivity of $\bowtie_{R1.x=R2.y}$ can be estimated as...

- if both x and y are the keys: $\frac{1}{max(|R1|,|R2|)}$
- if only x is the key:

Homework: Task 2

We know $|R1|$, $|R2|$, $|R1.x|$, $|R2.y|$, and whether x and y are keys or not.
First, the size of $R1 \times R2$ is $|R1||R2|$
The selectivity of $\bowtie_{R1.x=R2.y}$ can be estimated as...

- if both x and y are the keys: $\frac{1}{max(|R1|,|R2|)}$
- if only x is the key: $\frac{1}{|R1|}$

Homework: Task 2

We know $|R1|$, $|R2|$, $|R1.x|$, $|R2.y|$, and whether x and y are keys or not.
First, the size of $R1 \times R2$ is $|R1||R2|$
The selectivity of $\bowtie_{R1.x=R2.y}$ can be estimated as...

- if both x and y are the keys: $\frac{1}{max(|R1|,|R2|)}$
- if only x is the key: $\frac{1}{|R1|}$
- if both x and y are not the keys:

Homework: Task 2

We know $|R1|$, $|R2|$, $|R1.x|$, $|R2.y|$, and whether x and y are keys or not.
First, the size of $R1 \times R2$ is $|R1||R2|$
The selectivity of $\bowtie_{R1.x=R2.y}$ can be estimated as...

- if both x and y are the keys: $\frac{1}{max(|R1|,|R2|)}$
- if only x is the key: $\frac{1}{|R1|}$
- if both x and y are not the keys: $\frac{1}{max(|R1.x|,|R2.y|)}$

Selectivity estimation

We know $|R1|$, $max(R1.x)$, $min(R1.x)$, $R1.x$ is arithmetic.

The selectivity of $\sigma_{R1.x>c}$ is

Selectivity estimation

We know $|R1|$, $max(R1.x)$, $min(R1.x)$, $R1.x$ is arithmetic.

The selectivity of $\sigma_{R1.x>c}$ is $\frac{max(R1.x)-c}{max(R1.x)-min(R1.x)}$

Selectivity estimation

We know $|R1|$, $max(R1.x)$, $min(R1.x)$, $R1.x$ is arithmetic.

The selectivity of $\sigma_{R1.x>c}$ is $\frac{max(R1.x)-c}{max(R1.x)-min(R1.x)}$

The selectivity of $\sigma_{c1<R1.x<c2}$ is

Selectivity estimation

We know $|R1|$, $max(R1.x)$, $min(R1.x)$, $R1.x$ is arithmetic.

The selectivity of $\sigma_{R1.x>c}$ is $\frac{max(R1.x)-c}{max(R1.x)-min(R1.x)}$

The selectivity of $\sigma_{c1<R1.x<c2}$ is $\frac{c2-c1}{max-min}$

Homework: Task 3

- $|R| = 1,000$ pages, $|S| = 100,000$ pages
- 1 page - 50 tuples, 1 block - 100 pages
- avg. access = 10 ms, transfer speed = 10,000 pages/sec
- Time for block-nested loops join = ?

Homework: Task 3

- $|R| = 1,000$ pages, $|S| = 100,000$ pages
- 1 page - 50 tuples, 1 block - 100 pages
- avg. access = 10 ms, transfer speed = 10,000 pages/sec
- Time for block-nested loops join = ?
- choose left argument: R vs. S, $\frac{1,000}{100}$ vs. $\frac{100,000}{100} \Rightarrow R$

Homework: Task 3

- Time to read one block:
  $T_b = avg.seek + (100 \frac{1}{\text{transfer speed}}) = 0.02s$

Homework: Task 3

- Time to read one block:
  $T_b = avg.seek + (100 \frac{1}{\text{transfer speed}}) = 0.02s$
- Read 1 block from R, join it with S:

$$T_b + \text{time to read S} \approx 10s$$

Homework: Task 3

- Time to read one block:
  $T_b = avg.seek + (100 \frac{1}{\text{transfer speed}}) = 0.02s$
- Read 1 block from R, join it with S:

$$T_b + \text{time to read S} \approx 10s$$

- Repeat it for every block in R:

$$T_{BNLJ} = \frac{\#\text{pages in R}}{\text{block size}}(10s) \approx 100s$$

Query Graphs

```
select v.titel
  from Vorlesungen v, Professoren p
  where v.gelesenvon = p.persnr
    and p.name = 'Kant'
    and v.sws = 2;
```

Query Graphs

```
select r.a, s.c
  from R r, S s, T t, U u
  where r.a = s.a
    and r.b = t.b
    and r.b = u.b;
```

Query Graphs

```
select r.a, s.c
  from R r, S s
  where r.a + s.a = 7;
```

Query Graphs

```
select r.a, s.c
  from R r, S s, T t, U u
  where (r.a + s.b) = (t.b + u.a);
```

Search space

Search space is defined by:

- Query graph type

Search space

Search space is defined by:

- Query graph type (chain, star, tree, clique,cycle, grid)
- Join tree class

Search space

Search space is defined by:

- Query graph type (chain, star, tree, clique,cycle, grid)
- Join tree class (left-deep, zig-zag, bushy)
- Cost function class

Search space

```
select *
from R1, R2, R3, R4
where R1.a = R2.b
      and R2.c = R3.d
      and R3.e = R4.f
```

- What kind of query graph is it?

Search space

```
select *
from R1, R2, R3, R4
where R1.a = R2.b
      and R2.c = R3.d
      and R3.e = R4.f
```

- What kind of query graph is it?
- Let's allow cross-products $\Rightarrow$ the shape of the query graph does not matter

Search space

```
select *
from R1, R2, R3, R4
where R1.a = R2.b
      and R2.c = R3.d
      and R3.e = R4.f
```

- What kind of query graph is it?
- Let's allow cross-products $\Rightarrow$ the shape of the query graph does not matter
- Count left-deep trees
- Count zig-zag trees
- Count bushy trees

Homework: Task 1 (15 points)

Selectivity estimation continues...

- Our estimations (prev.homework) are far from perfect
- Construct two specific SQL examples, where our estimations are very "bad"
- Use IMDB schema (SQL script on the website)
- "Bad" – means that the logical plan will be suboptimal (w.r.t C_{out}), if we use these estimations
- In other words, bad estimations mislead the optimizer and it outputs a clearly suboptimal plan
- Construct one query where PostgreSQL estimates are wrong and lead to a suboptimal plan
- Force the optimal plan, compare the difference
- Hint: use join_collapse_limit
- Hint: correlations on joins, LIKE predicate etc.

Homework: Task 2 (5 points)

- Give an example query instance where the optimal join tree (using C_{out}) is bushy and includes a cross product.
- Note: the query graph should be connected!

Info

- Exercises due: 9 AM, November 10, 2014