



THE UNIVERSITY
of EDINBURGH

Advanced Database Systems

Spring 2024

Lecture #30:

Revision II

PLAN FOR TODAY

Selectivity Estimation

Query Optimisation

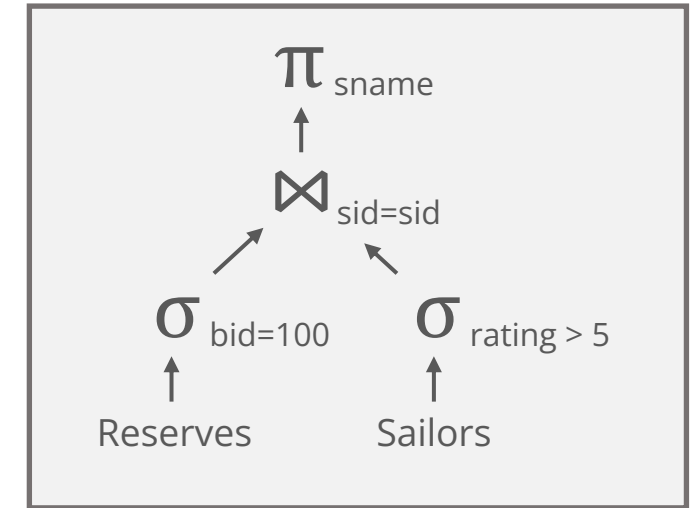
SELECTIVITY ESTIMATION

To estimate the cost of a query, we add up the estimated costs of each operator in the query

Need to know the size of the **intermediate relations** (generated from one operator and passed into another) in order to do this!

Need to know the **selectivity** of predicates - what % of tuples are selected by a predicate

These are all estimates... if we don't know, we make up a value for it (e.g., selectivity = 1/10)



SELECTIVITY ESTIMATION: EQUALITIES

PREDICATE	SELECTIVITY	ASSUMPTION
$A = \text{value}$	$1 / (\# \text{ distinct values of } A \text{ in relation})$	We know $ A $
$A = \text{value}$	$1 / 10$	We don't know $ A $
$A = B$	$1 / \text{MAX}(\# \text{ distinct } A\text{-values}, \# \text{ distinct } B\text{-values})$	We know $ A $ and $ B $
$A = B$	$1 / (\# \text{ distinct values of } A)$	We know $ A $ but not $ B $
$A = B$	$1 / 10$	We don't know $ A $ nor $ B $

$|\text{column}|$ = the number of distinct values for the column

If you have an index on column A, you can assume you know $|A|$, $\max(A)$, and $\min(A)$

When using selectivity to compute # of tuples, round up the result (e.g. $245.7 \rightarrow 246$ tuples)

SELECTIVITY ESTIMATION: INEQUALITIES ON INTEGERS

PREDICATE	SELECTIVITY	ASSUMPTION
$A < c$	$(c - \min(A)) / (\max(A) - \min(A) + 1)$	We know $\max(A)$ and $\min(A)$
$A \leq c$	$(c - \min(A) + 1) / (\max(A) - \min(A) + 1)$	c is an integer
$A < c$	$1 / 3$	We don't know $\max(A)$ and $\min(A)$
$A \leq c$		c is an integer
$A > c$	$(\max(A) - c) / (\max(A) - \min(A) + 1)$	We know $\max(A)$ and $\min(A)$
$A \geq c$	$(\max(A) - c + 1) / (\max(A) - \min(A) + 1)$	c is an integer
$A > c$	$1 / 3$	We don't know $\max(A)$ and $\min(A)$
$A \geq c$		c is an integer

* We add 1 to the denominator in order for our [low, high] range to be inclusive

E.g. range [2, 4] = 2, 3, 4 $\rightarrow (4 - 2) + 1 = 3$

SELECTIVITY ESTIMATION: INEQUALITIES ON FLOATS

PREDICATE	SELECTIVITY	ASSUMPTION
$A < c$ $A \leq c$	$(c - \min(A)) / (\max(A) - \min(A))$	We know $\max(A)$ and $\min(A)$ c is a float
$A < c$ $A \leq c$	$1 / 3$	We don't know $\max(A)$ and $\min(A)$ c is a float
$A > c$ $A \geq c$	$(\max(A) - c) / (\max(A) - \min(A))$	We know $\max(A)$ and $\min(A)$ c is a float
$A > c$ $A \geq c$	$1 / 3$	We don't know $\max(A)$ and $\min(A)$ c is a float

* We don't add 1 to the denominator (floats are continuous, integers are discrete)

E.g. range $[2.0, 4.0] = 2.0, 2.1, \dots, 3.9, 4.0 \rightarrow 4.0 - 2.0 = 2.0$

SELECTIVITY ESTIMATION: CONNECTIVES

PREDICATE	SELECTIVITY	ASSUMPTION
p1 AND p2	$\text{sel}(p1) \cdot \text{sel}(p2)$	Independent predicates
p1 OR p2	$\text{sel}(p1) + \text{sel}(p2) - \text{sel}(p1) \cdot \text{sel}(p2)$	Independent predicates
NOT p	$1 - \text{sel}(p)$	

SELECTIVITY ESTIMATION

How many tuples are selected by the following query?

```
SELECT * FROM R
```

R(A,B,C) has 1000 tuples

Attribute A:

50 unique integers, uniformly distributed in the range [1, 50]

Attribute B:

100 unique floats, uniformly distributed in the range [1, 100]

SELECTIVITY ESTIMATION

How many tuples are selected by the following query?

```
SELECT * FROM R
```

1000 tuples
(no predicates, select all)

R(A,B,C) has 1000 tuples

Attribute A:

50 unique integers, uniformly distributed in the range [1, 50]

Attribute B:

100 unique floats, uniformly distributed in the range [1, 100]

SELECTIVITY ESTIMATION

How many tuples are selected by the following query?

```
SELECT * FROM R
WHERE A = 42
```

R(A,B,C) has 1000 tuples

Attribute A:

50 unique integers, uniformly distributed in the range [1, 50]

Attribute B:

100 unique floats, uniformly distributed in the range [1, 100]

SELECTIVITY ESTIMATION

How many tuples are selected by the following query?

```
SELECT * FROM R
WHERE A = 42
```

50 unique values in A
 $1/50 \cdot (1000 \text{ tuples}) = 20 \text{ tuples}$

R(A,B,C) has 1000 tuples

Attribute A:

50 unique integers, uniformly distributed in the range [1, 50]

Attribute B:

100 unique floats, uniformly distributed in the range [1, 100]

SELECTIVITY ESTIMATION

How many tuples are selected by the following query?

```
SELECT * FROM R
WHERE C = 42
```

R(A,B,C) has 1000 tuples

Attribute A:

50 unique integers, uniformly distributed in the range [1, 50]

Attribute B:

100 unique floats, uniformly distributed in the range [1, 100]

SELECTIVITY ESTIMATION

How many tuples are selected by the following query?

```
SELECT * FROM R
WHERE C = 42
```

No information about C

$1/10 \cdot (1000 \text{ tuples}) = 100 \text{ tuples}$

R(A,B,C) has 1000 tuples

Attribute A:

50 unique integers, uniformly distributed in the range [1, 50]

Attribute B:

100 unique floats, uniformly distributed in the range [1, 100]

SELECTIVITY ESTIMATION

How many tuples are selected by the following query?

```
SELECT * FROM R
WHERE A <= 25
```

R(A,B,C) has 1000 tuples

Attribute A:

50 unique integers, uniformly distributed in the range [1, 50]

Attribute B:

100 unique floats, uniformly distributed in the range [1, 100]

SELECTIVITY ESTIMATION

How many tuples are selected by the following query?

```
SELECT * FROM R
WHERE A <= 25
```

$$\begin{aligned} \text{sel}(A \leq 25) &= \\ &= (25 - 1 + 1) / (50 - 1 + 1) \\ &= 1/2 \end{aligned}$$

$$1/2 \cdot (1000 \text{ tuples}) = 500 \text{ tuples}$$

R(A,B,C) has 1000 tuples

Attribute A:

50 unique integers, uniformly distributed in the range [1, 50]

Attribute B:

100 unique floats, uniformly distributed in the range [1, 100]

SELECTIVITY ESTIMATION

How many tuples are selected by the following query?

```
SELECT * FROM R
WHERE B <= 25
```

R(A,B,C) has 1000 tuples

Attribute A:

50 unique integers, uniformly distributed in the range [1, 50]

Attribute B:

100 unique floats, uniformly distributed in the range [1, 100]

SELECTIVITY ESTIMATION

How many tuples are selected by the following query?

```
SELECT * FROM R
WHERE B <= 25
```

$$\begin{aligned} \text{sel}(B \leq 25) &= \\ &= (25 - 1) / (100 - 1) \\ &= 24/99 = 0.2424\dots \end{aligned}$$

$$\text{round}(0.2424\dots \cdot (1000 \text{ tuples})) = 242 \text{ tuples}$$

R(A,B,C) has 1000 tuples

Attribute A:

50 unique integers, uniformly distributed in the range [1, 50]

Attribute B:

100 unique floats, uniformly distributed in the range [1, 100]

SELECTIVITY ESTIMATION

How many tuples are selected by the following query?

```
SELECT * FROM R
WHERE C <= 25
```

R(A,B,C) has 1000 tuples

Attribute A:

50 unique integers, uniformly distributed in the range [1, 50]

Attribute B:

100 unique floats, uniformly distributed in the range [1, 100]

SELECTIVITY ESTIMATION

How many tuples are selected by the following query?

```
SELECT * FROM R
WHERE C <= 25
```

No information about C
 $\text{round}(1/3 \cdot (1000 \text{ tuples})) = 333 \text{ tuples}$

R(A,B,C) has 1000 tuples

Attribute A:

50 unique integers, uniformly distributed in the range [1, 50]

Attribute B:

100 unique floats, uniformly distributed in the range [1, 100]

SELECTIVITY ESTIMATION

How many tuples are selected by the following query?

```
SELECT * FROM R
WHERE A <= 25
      AND B <= 25
```

R(A,B,C) has 1000 tuples

Attribute A:

50 unique integers, uniformly distributed in the range [1, 50]

Attribute B:

100 unique floats, uniformly distributed in the range [1, 100]

SELECTIVITY ESTIMATION

How many tuples are selected by the following query?

```
SELECT * FROM R
WHERE A <= 25
      AND B <= 25
```

$$\begin{aligned} & \text{sel}(A \leq 25) \cdot \text{sel}(B \leq 25) \\ &= 1/2 \cdot 24/99 = 12/99 \\ &= 0.1212\dots \end{aligned}$$

$$\text{round}(0.1212\dots \cdot (1000 \text{ tuples})) = 121 \text{ tuples}$$

R(A,B,C) has 1000 tuples

Attribute A:

50 unique integers, uniformly distributed in the range [1, 50]

Attribute B:

100 unique floats, uniformly distributed in the range [1, 100]

SELECTIVITY ESTIMATION

How many tuples are selected by the following query?

```
SELECT * FROM R
WHERE A <= 25
      OR B <= 25
```

R(A,B,C) has 1000 tuples

Attribute A:

50 unique integers, uniformly distributed in the range [1, 50]

Attribute B:

100 unique floats, uniformly distributed in the range [1, 100]

SELECTIVITY ESTIMATION

How many tuples are selected by the following query?

```
SELECT * FROM R
WHERE A <= 25
      OR B <= 25
```

$$\begin{aligned} & \text{sel}(A \leq 25) + \text{sel}(B \leq 25) \\ & - \text{sel}(A \leq 25) \cdot \text{sel}(B \leq 25) \\ & = 1/2 + 24/99 - 1/2 \cdot 24/99 = 0.62121\dots \end{aligned}$$

$$\text{round}(0.62121\dots \cdot (1000 \text{ tuples})) = 621 \text{ tuples}$$

R(A,B,C) has 1000 tuples

Attribute A:

50 unique integers, uniformly distributed in the range [1, 50]

Attribute B:

100 unique floats, uniformly distributed in the range [1, 100]

SELECTIVITY ESTIMATION

How many tuples are selected by the following query?

```
SELECT * FROM R
WHERE A = C
```

R(A,B,C) has 1000 tuples

Attribute A:

50 unique integers, uniformly distributed in the range [1, 50]

Attribute B:

100 unique floats, uniformly distributed in the range [1, 100]

SELECTIVITY ESTIMATION

How many tuples are selected by the following query?

```
SELECT * FROM R
WHERE A = C
```

No information about C
 $1/50 \cdot (1000 \text{ tuples}) = 20 \text{ tuples}$

R(A,B,C) has 1000 tuples

Attribute A:

50 unique integers, uniformly distributed in the range [1, 50]

Attribute B:

100 unique floats, uniformly distributed in the range [1, 100]

SELECTIVITY ESTIMATION

How many tuples are selected by the following query?

```
SELECT * FROM R, S
WHERE R.A = S.D
```

R(A,B,C) has 1000 tuples

S(D, E) has 500 tuples

Attribute R.A:

50 unique integers, uniformly distributed in the range [1, 50]

Attribute S.D:

25 unique integers, uniformly distributed in the range [1, 25]

SELECTIVITY ESTIMATION

How many tuples are selected by the following query?

```
SELECT * FROM R, S
WHERE R.A = S.D
```

Max output size = $|R| \cdot |S|$

$sel(R.A = S.D)$

$= 1 / \text{MAX}(50, 25) = 1/50$

$1/50 \cdot (1000 \cdot 500) = 10,000$ tuples

R(A,B,C) has 1000 tuples

S(D, E) has 500 tuples

Attribute R.A:

50 unique integers, uniformly distributed in the range [1, 50]

Attribute S.D:

25 unique integers, uniformly distributed in the range [1, 25]

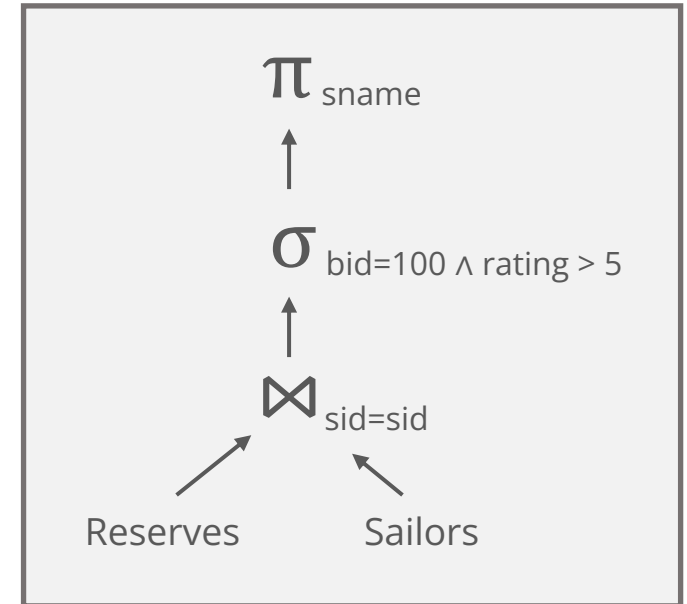
QUERY OPTIMISATION

QUERY OPTIMISATION – BACKGROUND

We can represent relational algebra expressions as trees

Order of operators affects I/Os and resource usage, but not necessarily output

$$\pi_{sname}(\sigma_{bid=100 \wedge rating > 5}(\text{Reserves} \bowtie \text{Sailors}))$$



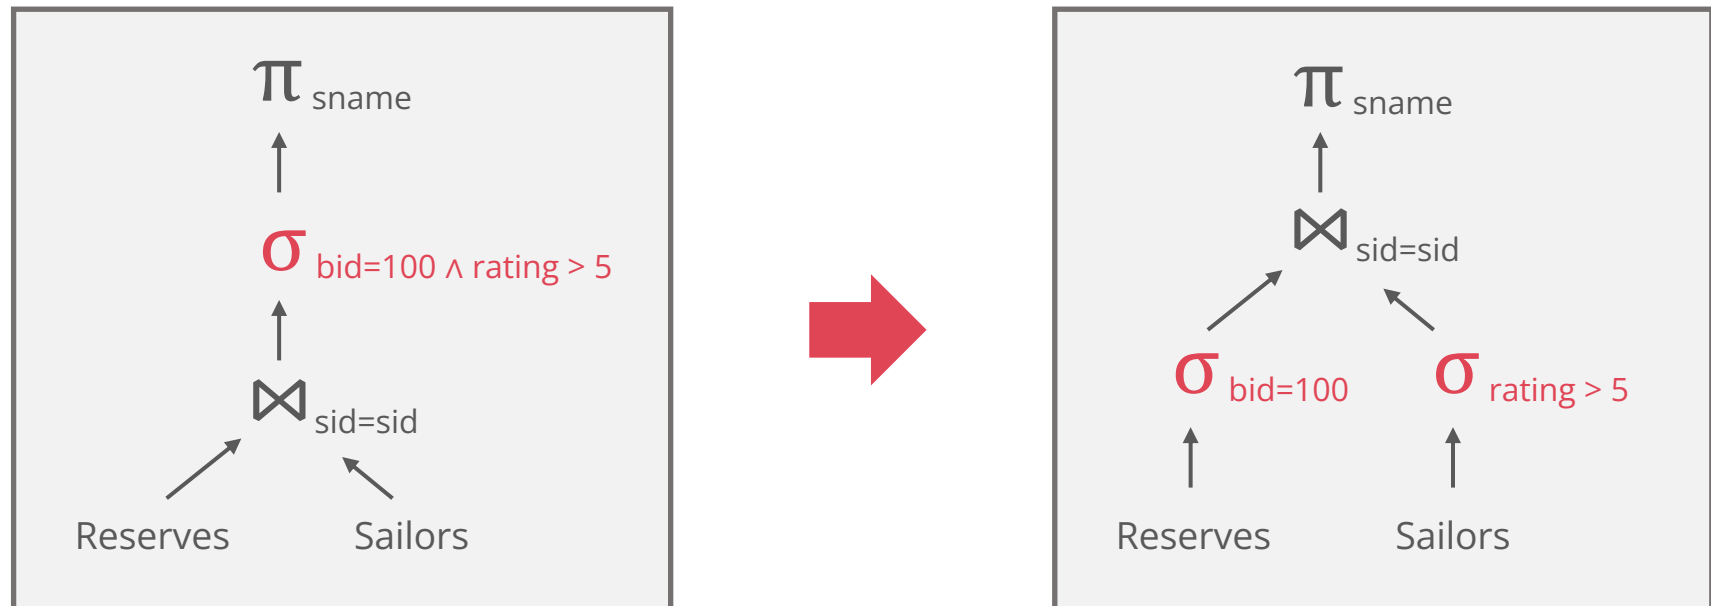
QUERY OPTIMISATION – ALTERNATE PLANS

Given a plan, some things we can do are:

Push selections/projections down the tree

The earlier we reduce the size of input, the fewer I/Os are incurred as we traverse up the tree

Reduces I/O cost if materialized



QUERY OPTIMISATION – ALTERNATE PLANS

Given a plan, some things we can do are:

- Push selections/projections down the tree

- Materialize intermediate relations (write to a temp file)

 - Results in additional write I/Os, but is better in the long run

- Use indices (e.g., INLJ)

QUERY OPTIMISATION - MATERIALISING

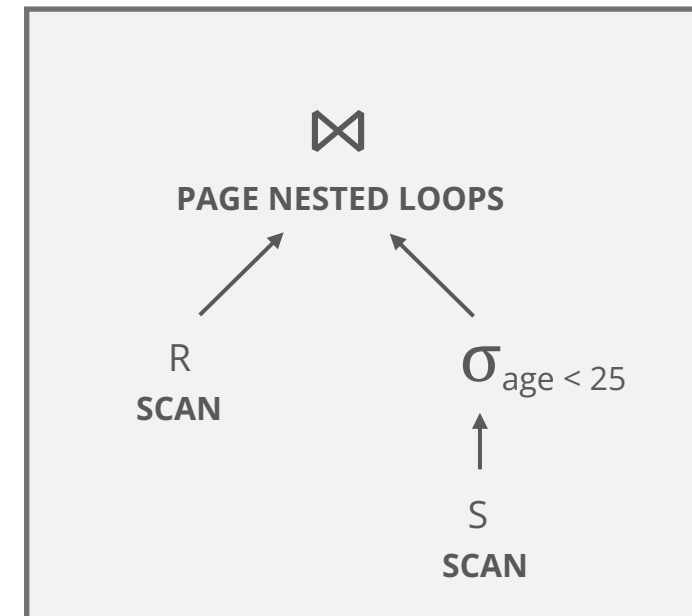
Table R consists of 50 pages

Table S consists of 100 pages

$\text{sel}(S.\text{age} < 25) = 0.5$

Without materializing, we are performing $\sigma_{\text{age} < 25}$ on the fly each time in PNLJ, and scanning the entire table S for each page of R

Cost = Scan R (50) + PNLJ ($50 \cdot 100$) = **5,050** I/Os



QUERY OPTIMISATION - MATERIALISING

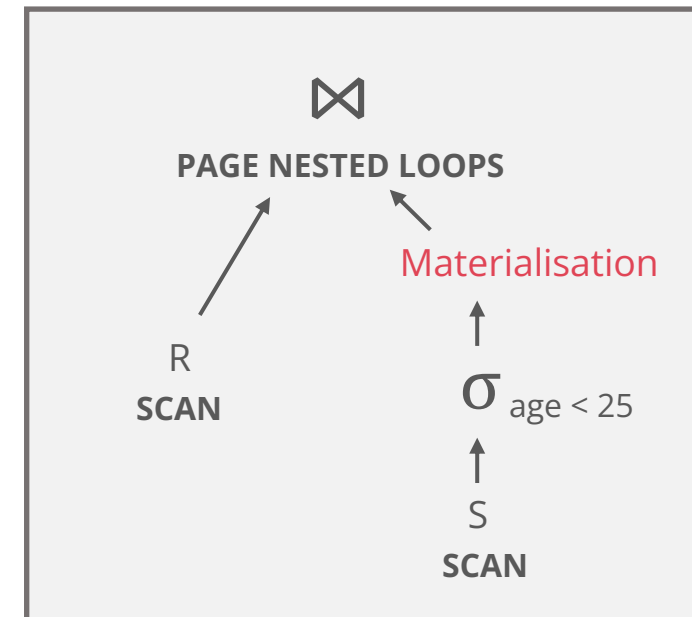
Table R consists of 50 pages

Table S consists of 100 pages

$\text{sel}(S.\text{age} < 25) = 0.5$

By materializing the intermediate relation, we are applying $\sigma_{\text{age} < 25}$ before PNLJ, and performing the join on the result of the selection

Cost = Scan R (50) + Scan S (100) +
 Materialise (50) + PNLJ ($50 \cdot 50$) = **2,700** I/Os



QUERY OPTIMISATION

A query optimiser takes in a query plan (e.g., one directly translated from a SQL query) and outputs a better (hopefully optimal) query plan

Works on and optimizes over a **plan space** (set of all plans considered)

Performs **cost estimation** on query plans

Uses a **search algorithm** to search through plan space to find plan with lowest cost estimate

May not be optimal (bad estimate or small plan space considered)

QUERY OPTIMISATION – SYSTEM R

We will be looking at the System R optimiser (aka Selinger optimiser)

Plan space

Only left-deep trees, avoid Cartesian products unless they are the only option

Left-deep trees represent a plan where all new tables are joined one at a time from the right

Cost estimation

Actual System R optimiser incorporates both CPU and I/O cost


We will only use I/O cost in this course

Search algorithm


Dynamic programming

QUERY OPTIMISATION – SYSTEM R

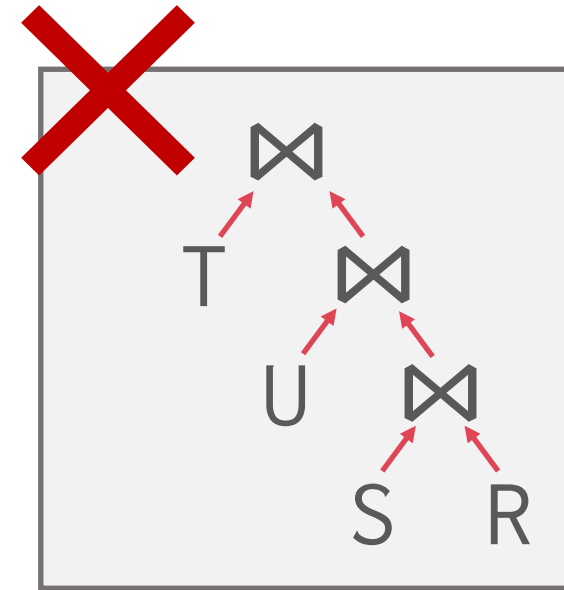
Only consider left-deep plans



left-deep
 $((R \bowtie S) \bowtie T) \bowtie U$



bushy
 $(R \bowtie S) \bowtie (T \bowtie U)$



right-deep
 $T \bowtie (U \bowtie (S \bowtie R))$

QUERY OPTIMISATION – SYSTEM R

Why only left-deep trees?

Join new tables one at a time from the right

Create an ordering in which to add tables to the query being executed

Too many possible trees for joins

Using only left-deep trees: $N!$ different ways to order relations

Including all permutations tree layouts: A very large number of ways to parenthesize given an ordering (superexponential in N)

# of relations n	# of different join trees
2	2
3	12
4	120
5	1,680
6	30,240
7	665,280
8	17,297,280
10	17,643,225,600

QUERY OPTIMISATION – SYSTEM R

Search algorithm for System R: use dynamic programming

Based on the principle of optimality

Runtime drops from $n!$ to around $n \cdot 2^n$

To be considered, must be:

Left deep

No Cartesian products

(I.e. if we join R and S on <cond1> and we join S and T on <cond2>, we don't consider joining R and T if there's no condition between them)

QUERY OPTIMISATION – SYSTEM R

For N relations joined, perform N passes

On the i -th pass, output only the best plan for joining any i of the N relations

Also keep around plans that have higher cost but have an **interesting order**

Interesting orders are orderings on intermediate relations that **may** help reduce the cost of later operators (e.g., joins, sorting, hashing)

ORDER BY attributes

GROUP BY attributes

downstream join attributes

SYSTEM R OPTIMISATION: EXAMPLE

Pass 1:

Find minimum cost access method for each (relation, interesting order) pair

Index scan, full table scans

A toy example:

```
SELECT * FROM R, S, T
WHERE R.B = S.B
      AND S.C = T.C
      AND R.A <= 50
```


SYSTEM R OPTIMISATION: EXAMPLE

Pass 1:

Assume the single table access plans have the following IO costs:

```
SELECT * FROM R, S, T
WHERE R.B = S.B
      AND S.C = T.C
      AND R.A <= 50
```

Full scan on R	1000 I/Os	
Index scan on R.A	200 I/Os	(sorted on R.A)
Index scan on R.B	1100 I/Os	(sorted on R.B)
Full scan on S	2000 I/Os	
Index scan on S.B	2500 I/Os	(sorted on S.B)
Full scan on T	3000 I/Os	
Index scan on T.C	3500 I/Os	(sorted on T.C)
Index scan on T.D	3500 I/Os	(sorted on T.D)

SYSTEM R OPTIMISATION: EXAMPLE

Pass 1:

Which single table access plans advance to the next stage?

```
SELECT * FROM R, S, T
WHERE R.B = S.B
      AND S.C = T.C
      AND R.A <= 50
```

Full scan on R	1000 I/Os	
Index scan on R.A	200 I/Os	(sorted on R.A)
Index scan on R.B	1100 I/Os	(sorted on R.B)
Full scan on S	2000 I/Os	
Index scan on S.B	2500 I/Os	(sorted on S.B)
Full scan on T	3000 I/Os	
Index scan on T.C	3500 I/Os	(sorted on T.C)
Index scan on T.D	3500 I/Os	(sorted on T.D)

SYSTEM R OPTIMISATION: EXAMPLE

Pass 1:

Which single table access plans advance to the next stage?

```
SELECT * FROM R, S, T
WHERE R.B = S.B
      AND S.C = T.C
      AND R.A <= 50
```

Full scan on R	1000 I/Os	
Index scan on R.A	200 I/Os	(sorted on R.A)
Index scan on R.B	1100 I/Os	(sorted on R.B)
Full scan on S	2000 I/Os	
Index scan on S.B	2500 I/Os	(sorted on S.B)
Full scan on T	3000 I/Os	
Index scan on T.C	3500 I/Os	(sorted on T.C)
Index scan on T.D	3500 I/Os	(sorted on T.D)

SYSTEM R OPTIMISATION: EXAMPLE

Pass 1:

Which single table access plans advance to the next stage?

Index scan on R.A (best overall plan for R)

Index scan on R.B (output sorted on R.B and R.B is used in a join)

Full scan on S (best overall plan for S)

Index scan on S.B (output sorted on S.B and S.B is used in a join)

Full scan on T (best overall plan for T)

Index scan on T.C (output sorted on T.C and T.C is used in a join)

SYSTEM R OPTIMISATION: EXAMPLE

Assume the following join costs:

$R \bowtie_{\text{BNLJ}} S$	21,000 I/Os
$R \bowtie_{\text{SMJ}} S$	3,600 I/Os
$S \bowtie_{\text{BNLJ}} R$	18,000 I/Os
$S \bowtie_{\text{SMJ}} R$	3,000 I/Os

$R \bowtie_{\text{BNLJ}} T$	30,000 I/Os
$R \bowtie_{\text{SMJ}} T$	40,000 I/Os
$T \bowtie_{\text{BNLJ}} R$	35,000 I/Os
$T \bowtie_{\text{SMJ}} R$	20,000 I/Os

$S \bowtie_{\text{BNLJ}} T$	15,000 I/Os
$S \bowtie_{\text{SMJ}} T$	10,000 I/Os
$T \bowtie_{\text{BNLJ}} S$	25,000 I/Os
$T \bowtie_{\text{SMJ}} S$	30,000 I/Os

```
SELECT * FROM R, S, T
WHERE R.B = S.B
      AND S.C = T.C
      AND R.A <= 50
```

Which of these joins will actually be considered by the query optimiser on pass 2?

SYSTEM R OPTIMISATION: EXAMPLE

Assume the following join costs:

$R \bowtie_{\text{BNLJ}} S$ 21,000 I/Os

$R \bowtie_{\text{SMJ}} S$ 3,600 I/Os

$S \bowtie_{\text{BNLJ}} R$ 18,000 I/Os

$S \bowtie_{\text{SMJ}} R$ 3,000 I/Os

$R \bowtie_{\text{BNLJ}} T$ 30,000 I/Os

$R \bowtie_{\text{SMJ}} T$ 40,000 I/Os

$T \bowtie_{\text{BNLJ}} R$ 35,000 I/Os

$T \bowtie_{\text{SMJ}} R$ 20,000 I/Os

$S \bowtie_{\text{BNLJ}} T$ 15,000 I/Os

$S \bowtie_{\text{SMJ}} T$ 10,000 I/Os

$T \bowtie_{\text{BNLJ}} S$ 25,000 I/Os

$T \bowtie_{\text{SMJ}} S$ 30,000 I/Os

```
SELECT * FROM R, S, T
WHERE R.B = S.B
      AND S.C = T.C
      AND R.A <= 50
```

Which of these joins will actually be considered by the query optimiser on Pass 2?

Eliminate Cartesian products (joins between R and T)

SYSTEM R OPTIMISATION: EXAMPLE

Assume the following join costs:

$R \bowtie_{\text{BNLJ}} S$	21,000 I/Os
$R \bowtie_{\text{SMJ}} S$	3,600 I/Os
$S \bowtie_{\text{BNLJ}} R$	18,000 I/Os
$S \bowtie_{\text{SMJ}} R$	3,000 I/Os

$R \bowtie_{\text{BNLJ}} T$	30,000 I/Os
$R \bowtie_{\text{SMJ}} T$	40,000 I/Os
$T \bowtie_{\text{BNLJ}} R$	35,000 I/Os
$T \bowtie_{\text{SMJ}} R$	20,000 I/Os

$S \bowtie_{\text{BNLJ}} T$	15,000 I/Os
$S \bowtie_{\text{SMJ}} T$	10,000 I/Os
$T \bowtie_{\text{BNLJ}} S$	25,000 I/Os
$T \bowtie_{\text{SMJ}} S$	30,000 I/Os

```

SELECT * FROM R, S, T
WHERE R.B = S.B
      AND S.C = T.C
      AND R.A <= 50
  
```

Which of these joins will advance to the next pass of the query optimiser?

SYSTEM R OPTIMISATION: EXAMPLE

Assume the following join costs:

R \bowtie_{BNLJ} S	21,000 I/Os
R \bowtie_{SMJ} S	3,600 I/Os
S \bowtie_{BNLJ} R	18,000 I/Os
S \bowtie_{SMJ} R	3,000 I/Os

R \bowtie_{BNLJ} T	30,000 I/Os
R \bowtie_{SMJ} T	40,000 I/Os
T \bowtie_{BNLJ} R	35,000 I/Os
T \bowtie_{SMJ} R	20,000 I/Os

S \bowtie_{BNLJ} T	15,000 I/Os
S \bowtie_{SMJ} T	10,000 I/Os
T \bowtie_{BNLJ} S	25,000 I/Os
T \bowtie_{SMJ} S	30,000 I/Os

```
SELECT * FROM R, S, T
WHERE R.B = S.B
      AND S.C = T.C
      AND R.A <= 50
```

Which of these joins will advance to the next pass of the query optimiser?

None of the joins produce an interesting order (no downstream joins, ORDER BY, GROUP BY). Only consider best join for each considered set of tables

SYSTEM R OPTIMISATION: EXAMPLE

Assume the following join costs:

R \bowtie_{BNLJ} S	21,000 I/Os
R \bowtie_{SMJ} S	3,600 I/Os
S \bowtie_{BNLJ} R	18,000 I/Os
S \bowtie_{SMJ} R	3,000 I/Os

R \bowtie_{BNLJ} T	30,000 I/Os
R \bowtie_{SMJ} T	40,000 I/Os
T \bowtie_{BNLJ} R	35,000 I/Os
T \bowtie_{SMJ} R	20,000 I/Os

S \bowtie_{BNLJ} T	15,000 I/Os
S \bowtie_{SMJ} T	10,000 I/Os
T \bowtie_{BNLJ} S	25,000 I/Os
T \bowtie_{SMJ} S	30,000 I/Os

```
SELECT * FROM R, S, T
WHERE R.B = S.B
      AND S.C = T.C
      AND R.A <= 50
```

Will any of these remaining joins produce an interesting order?

SYSTEM R OPTIMISATION: EXAMPLE

Assume the following join costs:

$R \bowtie_{\text{BNLJ}} S$	21,000 I/Os
$R \bowtie_{\text{SMJ}} S$	3,600 I/Os
$S \bowtie_{\text{BNLJ}} R$	18,000 I/Os
$S \bowtie_{\text{SMJ}} R$	3,000 I/Os

$R \bowtie_{\text{BNLJ}} T$	30,000 I/Os
$R \bowtie_{\text{SMJ}} T$	40,000 I/Os
$T \bowtie_{\text{BNLJ}} R$	35,000 I/Os
$T \bowtie_{\text{SMJ}} R$	20,000 I/Os

$S \bowtie_{\text{BNLJ}} T$	15,000 I/Os
$S \bowtie_{\text{SMJ}} T$	10,000 I/Os
$T \bowtie_{\text{BNLJ}} S$	25,000 I/Os
$T \bowtie_{\text{SMJ}} S$	30,000 I/Os

```
SELECT * FROM R, S, T
WHERE R.B = S.B
      AND S.C = T.C
      AND R.A <= 50
```

Will any of these remaining joins produce an interesting order?

No. $R \bowtie_{\text{SMJ}} S$ is sorted on B (not interesting), and $T \bowtie_{\text{SMJ}} S$ is sorted on C (not interesting)

SYSTEM R OPTIMISATION: EXAMPLE

How could we modify the query so that $S \bowtie_{SMJ} R$ yields an interesting order?

SYSTEM R OPTIMISATION: EXAMPLE

How could we modify the query so that $S \bowtie_{SMJ} R$ yields an interesting order?

$S \bowtie_{SMJ} R$ will be sorted on column B, so we need B to be interesting. We could add ORDER BY B, GROUP BY B, or another join condition involving R.B or S.B to the query to make it interesting

SYSTEM R OPTIMISATION: EXAMPLE

Will the query plan $T \bowtie_{\text{BNLJ}} (S \bowtie_{\text{SMJ}} R)$ be considered by the final pass of the query optimiser?

SYSTEM R OPTIMISATION: EXAMPLE

Will the query plan $T \bowtie_{\text{BNLJ}} (S \bowtie_{\text{SMJ}} R)$ be considered by the final pass of the query optimiser?

No, this query plan is not left-deep (all join results must be on the left side of their parent join), so it is not considered in the final pass

COMMIT
END TRANSACTION