

### Final Exam of Advanced Databases

#### Exercise 01 (3 marks)

A block nested loop join is executed with S as the outer relation and R as the inner relation. The join completes in 2 iterations. In each iteration, a block of 3 pages of S are loaded into the buffer. For each block of S, 10 pages of R are loaded into the buffer one by one

- 1) How many pages does S have, and how many pages does R have? Justify your answer.
- 2) Compute the total number of pages I/Os based on the simulation steps.

1)  $S = 6$  pages,  $R = 10$  pages

2) Iteration 1: load 3 pages of S  $\rightarrow 3$  I/Os for S

For each S block, load R (10 pages)  $\rightarrow 10$  I/Os for R

Iteration 2: load next 3 pages of S  $\rightarrow 3$  I/Os for S

For each S block, load R  $\rightarrow 10$  I/Os for R

Total  $6+2*10 = 26$  IOs

#### Exercise 02 (3 marks)

Given the Sort-Merge Join algorithm, where during sorting phase we observed:

- After the initial sort pass of R, the system produces 2 sorted runs of 5 pages each. Then, these 2 runs are merged in one merge pass
- After the initial sort pass of S, the system produces 2 sorted runs of 5 and 1 pages. Then, these 2 runs are merged in one merge pass

1) Deduce the buffer size and number of merge passes for each relation during the sorting phase.

2) Compute the total number of pages I/Os for  $R \bowtie S$  using Sort-Merge Join algorithm.

1) Buffer size (B)

In external sorting, the initial run size = B pages (all fit in memory).

- For R, initial runs are 5 pages  $\Rightarrow B = 5$  pages
- For S, one run is 5 pages (and the remainder is 1 page)  $\Rightarrow$  consistent with  $B = 5$  pages

So Buffer size = 5 pages

Number of merge passes

R: 2 initial runs → merged into 1 run in 1 merge pass

S: 2 initial runs → merged into 1 run in 1 merge pass

2) Total page I/Os for R $\bowtie$ SR  $\backslash$ bowtie SR $\bowtie$ S using Sort-Merge Join

- **Total for R sorting** read 10 + write 10 + read 10 + write 10 = 40 I/Os  
Initial sort pass      Merge pass
- **Total for S sorting** = read 6 + write 6 + read 6 + write 6 = 24 I/Os  
Initial sort pass      Merge pass

**Total Sort Merge Join cost = 40 (sort R)+24 (sort S)+10+6 (join)=80 page I/Os**

### Exercise 03 (4 marks)

Consider a B+ Tree of order d=3 used in the course (root and intermediate nodes: 5 pointers and 4 keys. Leaf node: 4 entries). The B+ Tree is initially empty.

1) Show the resulting tree after inserting the following keys into the B+ Tree in the sequence: 50, 30, 90, 100, 20, 10, 40, 60, 70, 80, 110

2) Now, starting from the tree obtained in Question 1, insert the following keys: 65, 66, 67, 75, 76, 77

[30 | 50 | 90]

/      |      |      \

[10,20, , ] [30,40, , ] [50,60,70,80] [90,100,110, ]

[65]

/      \

[30 | 50]      [70 | 76 | 90]

/ | \      / | | \

[10,20] [30,40] [50,60] [65,66,67] [70,75] [76,77,80] [90,100,110]

### Exercise 04 (10 marks)

Consider the following relational schema:

$\text{Suppliers}(\underline{\text{sid: integer}}, \text{pname: char(20)}, \text{city: char(20)})$   
 $\text{Supply}(\underline{\text{sid: integer}}, \underline{\text{pid: integer}})$   
 $\text{Parts}(\underline{\text{pid: integer}}, \text{pname: char(20)}, \text{price: real})$

And assuming

- suppliers are 11500 tuples stored on 500 pages
- supply has 192000 tuples stored on 1500 pages
- parts table has 22400 tuples stored on 700 pages
- 35% of parts have a price less than 1000 ( $\text{SF}(\text{prices} < 1000) = 0.35$ )
- 7% of suppliers from the city 'ELO' ( $\text{SF}(\text{city} = 'ELO') = 0.07$ )
- Assume the size of column types are: integer 4 bytes, char 1 byte, real 8 bytes
- The page size is 1024 bytes

Query1:  $\text{SELECT S.sid, city}$   
 $\text{FROM Supply Y, Suppliers S}$   
 $\text{WHERE Y.sid=S.sid AND S.city='ELO'}$

Query1:  $\text{SELECT pname, price}$   
 $\text{FROM Supply Y, Parts P}$   
 $\text{WHERE Y.pid=P.pid AND P.prices < 1000}$

Question) For the given queries:

- List all possible execution plans.
- For each plan, estimate the size (in pages) of the intermediate results at every required level of the plan.
- Estimate the total I/O cost of each plan using one join algorithm of your choice.

Page size = 1024 bytes

- Sizes:  
 $\text{sid} = 4 \text{ B}, \text{city} = 20 \text{ B} \rightarrow \text{supplier tuple} = 24 \text{ B}$   
 $\text{pname} = 20 \text{ B}, \text{price} = 8 \text{ B} \rightarrow \text{part tuple} = 28 \text{ B}$
- Selection factors:
  - $\text{city} = 'ELO' \rightarrow \text{SF} = 0.07$
  - $\text{price} < 1000 \rightarrow \text{SF} = 0.35$

Query 1 has 01 Plan Selection first, then join, selection in left side

$\Pi_{\text{sid}, \text{city}}(\Pi_{\text{sid}, \text{city}}(\sigma(\text{city} = 'ELO')(\text{Suppliers}))) \bowtie (\Pi_{\text{sid}}(\text{Supply}))$

## Estimation of the intermediate result sizes

$\Pi_{sid, city}(\Pi_{sid, city}(\sigma(city='ELO')(Suppliers)))$

- Suppliers filtered: Tuples=11500\*0.07=805 tuples
- Size per tuple = sid + city = 4 + 20 = 24 B
- Tuples per page = floor(1024 / 24) = 42 tuples per page
- Pages for filtered suppliers = ceil(805 / 42)  $\approx$  20 pages

Supply

- Supply = 192,000 tuples
- Size per tuple = sid = 4 B
- Tuples per page = floor(1024 / 4) = 256 tuples per page
- Pages for filtered suppliers = ceil(192000 / 256)  $\approx$  750 pages

Join cost using Block Nested Loop Join

$$500 + 20/(B-2) * 750$$

Join cost using Block Nested Loop Join

$$\text{Sort} \Pi_{sid, city}(\Pi_{sid, city}(\sigma(city='ELO')(Suppliers))) = 500 + 20 + \log_{B-1} 20/B$$

$$\text{Sort}(\Pi_{sid}(Supply)) = 1500 + 750 + \log_{B-1} 750/B$$

$$\text{Sort} \Pi_{sid, city}(\Pi_{sid, city}(\sigma(city='ELO')(Suppliers))) + \text{Sort}(\Pi_{sid}(Supply)) + [\Pi_{sid, city}(\Pi_{sid, city}(\sigma(city='ELO')(Suppliers)))] + [(\Pi_{sid}(Supply))] = 3540 + \log_{B-1} 20/B + \log_{B-1} 750/B$$

$$500 + 20 + \log_{B-1} \frac{20}{B} + 1500 + 750 + \log_{B-1} \frac{750}{B} + 20 + 750$$

$$\text{Total} = 3540 + \log_{B-1} \frac{20}{B} + \log_{B-1} \frac{750}{B}$$

$$3540 + \log_{B-1} \frac{15000}{B^2}$$

**Query 1 has 02 Plan Selection first, then join, selection in left side**

$\Pi_{sid, city(\text{Prname}, \text{price} (\sigma(\text{prices} < 1000')(\text{Parts})))} \bowtie (\Pi_{pid}(\text{Supply}))$

### Estimation of the intermediate result sizes

$\Pi_{sid, city(\text{Prname}, \text{price} (\sigma(\text{prices} < 1000')(\text{Parts})))}$

- Suppliers filtered: Tuples=22400\*0.35=7840 tuples
- Size per tuple = pname+ price= 20 + 8 = 28 B
- Tuples per page = floor(1024 / 28) = 36 tuples per page
- Pages for filtered suppliers = ceil(7840 / 36)  $\approx$  217 pages

Supply

- Supply = 192,000 tuples
- Size per tuple = pid = 4B
- Tuples per page = floor(1024/4) = 256 tuples per page
- Pages for filtered suppliers = ceil(192000 / 256)  $\approx$  750 pages

### Join cost using Block Nested Loop Join

$700 + 217/(B-2) * 750$

### Join cost using Block Nested Loop Join

$\Pi_{sid, city(\text{Prname}, \text{price} (\sigma(\text{prices} < 1000')(\text{Parts})))} = 700 + 217 + \log_{B-1} 217/B$

$\text{Sort}(\Pi_{pid}(\text{Supply})) = 1500 + 750 + \log_{B-1} 750/B$

$\Pi_{sid, city(\text{Prname}, \text{price} (\sigma(\text{prices} < 1000')(\text{Parts})))} + \text{Sort}(\Pi_{pid}(\text{Supply})) + [\Pi_{sid, city(\text{Prname}, \text{price} (\sigma(\text{prices} < 1000')(\text{Parts})))}] + [(\Pi_{pid}(\text{Supply}))] =$

$$700 + 217 + \log_{B-1} \frac{217}{B} + 1500 + 750 + \log_{B-1} \frac{750}{B} + 217 + 750$$

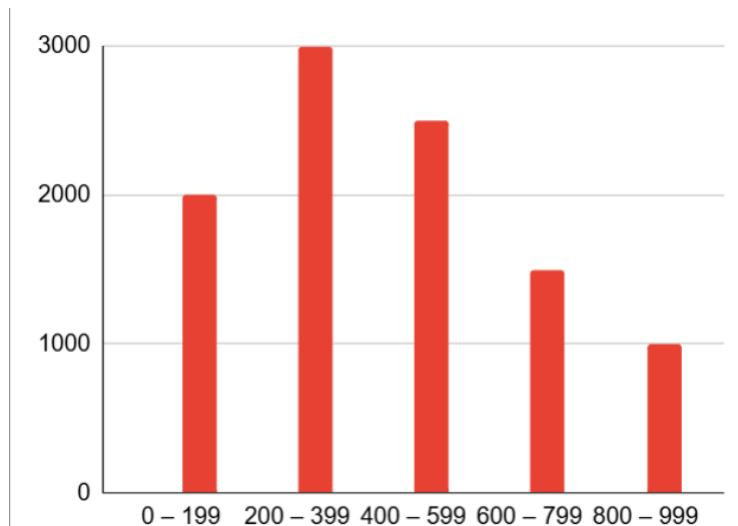
$$4134 + \log_{B-1} \frac{162750}{B^2}$$

**NOTE: All other correct solutions are considered**

**Exercise 05 (3 marks, BONUS)**

The Orders table has 10,000 rows, and we want to estimate the selectivity of a query condition using a histogram. Attribute `order_amount` has the following equally-width histogram with 5 buckets:

Bucket	Range	Frequency
1	0 – 199	2000
2	200 – 399	3000
3	400 – 599	2500
4	600 – 799	1500
5	800 – 999	1000



Estimate the selectivity factor for the following conditions:

1. `order_amount` = 350
2. `order_amount` < 500
3. `order_amount` BETWEEN 300 AND 700

### 1) `order_amount` = 350

350 in bucket 2 200 – 399 3000

bucket width=200

$$\text{Frequency per value} = \frac{3000}{200} = 15$$

$$\text{Selectivity factor} = \frac{15}{10000} = 0.0015$$

### 2) `order_amount` < 500

Bucket 1 (0–199) → fully included: 2000

Bucket 2 (200–399) → fully included: 3000

Bucket 3 (400–599) → partially included: 400–499 (100 values out of 200) → frequency proportion =  $2500 \times (100/200) = 1250$

Total rows=2000+3000+1250=6250

Selectivity factor=6250/10000=0.625

### 3) **order\_amount BETWEEN 300 AND 700**

Bucket 2 (200–399) → overlap 300–399 (100 values / 200) → frequency =  $3000 \times (100/200) = 1500$

Bucket 3 (400–599) → fully included → frequency = 2500

Bucket 4 (600–799) → overlap 600–700 (101 values / 200) → frequency  $\approx 1500 \times (101/200) \approx 757.5 \approx 758$

Total rows=1500+2500+758=4758

Selectivity factor=100004758≈0.476