

# CS143: Evaluation and Optimization

Ch 13-14

1

## Basic Steps in Query Processing

### 1. Parsing and translation:

... Translation of SQL queries into naive RA

### 2. Optimization: of RA expressions by

1. Pushing selection and projection
2. Estimating cost of different join orders, and selecting that of minimum cost
3. Selecting best implementation for each operator.

### 3. Evaluation

2

## Join Operation

- Several different algorithms to implement joins
  - Nested-loop join
  - Block nested-loop join
  - Indexed nested-loop join
  - Merge-join
  - Hash-join
- Choice based on cost estimate

3

## R JOIN S?

### Nested Join

	C			C	
R	40	T1	S	10	T6
	60	T2		60	T7
	30	T3		40	T8
	10	T4		20	T9
	20	T5			

4

- Nested-Loop Join
  - For each r ? R do
  - For each s ? S do
  - if r.C = s.C then output r,s pair

R		S	
40	T1	10	T6
60	T2	60	T7
30	T3	40	T8
10	T4	20	T9
20	T5		

R is the outer relation and S is the inner relation

5

## Block Nested-Loop Join

- Variant of nested-loop join in which every block of inner relation is paired with every block of outer relation.

```
for each block Br of r do begin
  for each block Bs of s do begin
    for each tuple tr in Br do begin
      for each tuple ts in Bs do begin
        Check if (tr, ts) satisfy the join condition
        if they do, add tr • ts to the result.
      end
    end
  end
end
end
```

6

- Index Join (conceptual)
  - If index supporting S.C does not exist , create one, and
  - For each r ? R do
    - X ? index-lookup(S.C, r.C)
    - For each s ? X, output (r,s)

C	T
40	T1
60	T2
30	T3
10	T4
20	T5

C	T
10	T6
60	T7
40	T8
20	T9

### Selection and Indexing

- If we have a S.C condition supported by an existing index we use the index
- If we have a conjunction, such as S.A>5 and S.B <10 with indexes on both, then we can select the better of the two (optimization)
- If there is no index, do we create one?

### Cost of Joins

- Nested Loop:**  $b_R + |R| \times b_S$  ( $b_R$  and  $b_S$  denote the number of blocks in R and S, resp.) this is the worst case. What is the best case?
- Pre-existing Index on S:**  $b_R + |R| \times k$  (k: depth of index)
- Block-Nested-Loop Join:**  $b_R + b_R \times b_S$  ( $b_S$ : number of blocks in S)

### Sort-Merge Join

- Sort the relations first and join

R	
10	T4
20	T5
30	T3
40	T1
60	T2

S	
10	T6
20	T9
40	T8
60	T7

### Merge-Join

- Can be used only for equi-joins and natural joins
- Sort both relations on their join attribute (if not already sorted on the join attributes).
- Merge the sorted relations to join them

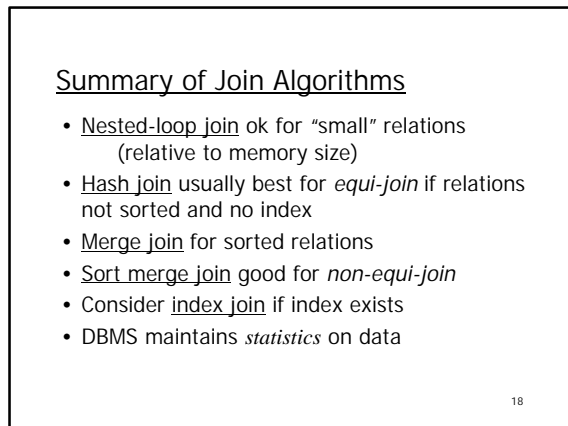
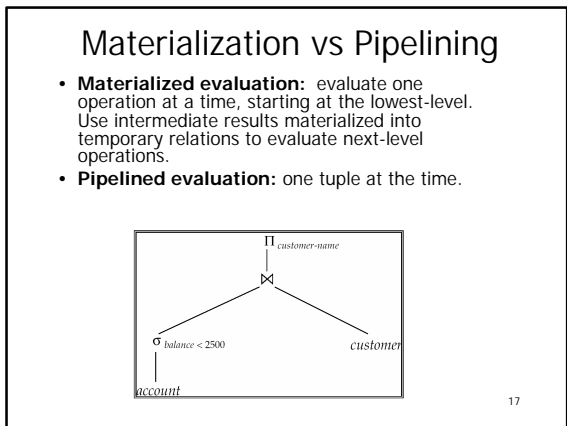
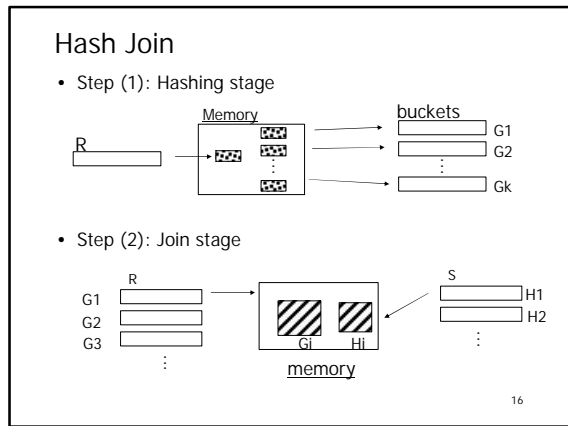
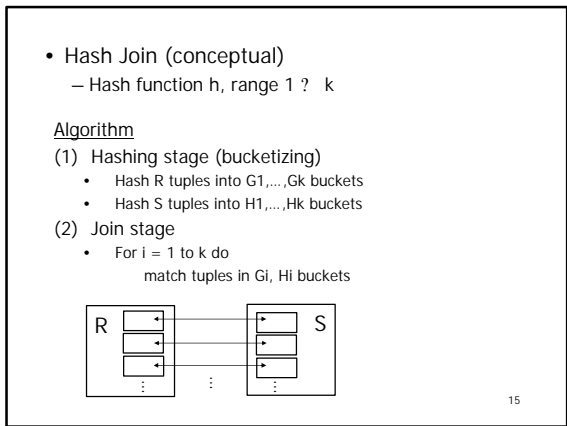
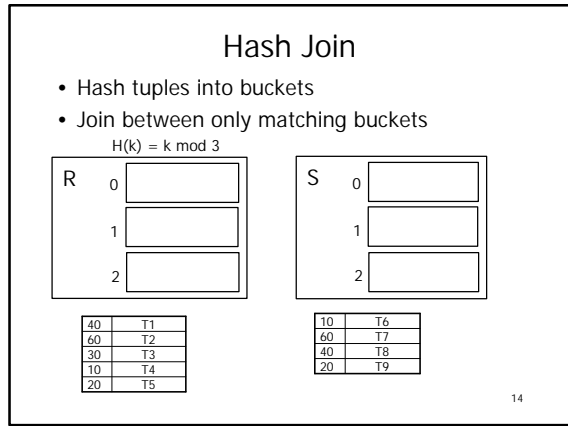
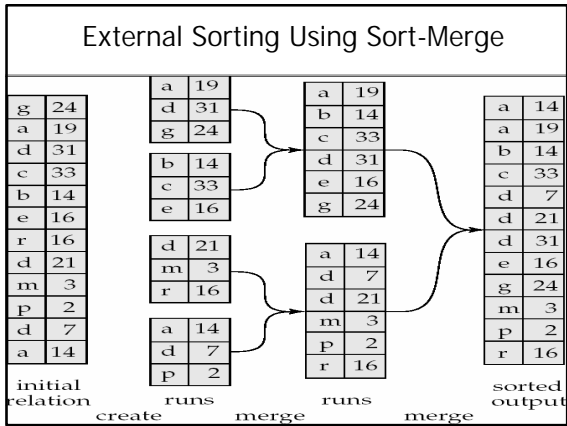
a1	a2
a	3
b	1
d	8
d	13
f	7
m	5
q	6

a1	a3
a	A
b	G
c	L
d	N
m	B

- Sort-Merge Join (conceptual)
  - if R and S not sorted, sort them
  - $i ? 1; j ? 1;$ 
    - While ( $i ? |R|$ ) ? ( $j ? |S|$ ) do
      - if  $R[i].C = S[j].C$  then outputTuples
      - else if  $R[i].C > S[j].C$  then  $j ? j+1$
      - else if  $R[i].C < S[j].C$  then  $i ? i+1$

R	
10	T4
20	T5
30	T3
40	T1
60	T2

S	
10	T6
20	T9
40	T8
60	T7



## Statistics collection commands

- DBMS has to collect statistics on tables/indexes
  - For optimal performance
  - Without stats, DBMS does stupid things...
- DB2
  - RUNSTATS ON TABLE <userid>.<table> AND INDEXES ALL
- Oracle
  - ANALYZE TABLE <table> COMPUTE STATISTICS
  - ANALYZE TABLE <table> ESTIMATE STATISTICS (cheaper than COMPUTE)
- Run the command after major update/index construction

19