

Cost Estimation of Relational Operations

- Cost estimation is important to improve the performance of relational operations
- There are two primary factors affecting the performance of relational operations:
 - Computation costs in CPU
 - Disk access costs
- Commonly, the disk access costs are dominant
 - 'ms' order
 - seek + rotational latency + data transfer
 - During a disk access, millions steps of a program can be executed in a recent CPU

2019/7/11

Advance Data Engineering (©H.Yokota)

121

Notations for the Cost Estimation

- We estimate the CPU costs by calculating the number of tuples to be handled and the disk access costs by calculating the number of disk pages to be accessed
- We use two notations
 - $\{R\}$: The number of tuples in the relation R (Cardinality)
 - To measure CPU costs
 - $|R|$: The number of disk pages for storing the relation R
 - To measure disk access costs

2019/7/11

Advance Data Engineering (©H.Yokota)

122

Relational Algebra Operations

- Dedicated Relational Algebra Operations
 - Selection ($\sigma_{\text{conditions}} R$)
 - Projection ($\pi_{\text{attribute-list}} R$)
 - Join (θ -Join: $R \bowtie_{\theta} S$, eq-Join: $R \bowtie_{r=s} S$)
 - Division ($R \div S$)
- Set Operations
 - Union ($R \cup S$)
 - Intersection ($R \cap S$)
 - Difference ($R - S$)
 - Cartesian Product ($R \times S$)

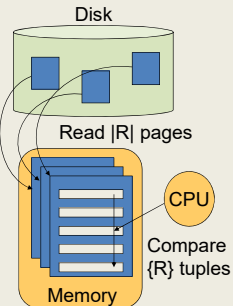
2019/7/11

Advance Data Engineering (©H.Yokota)

123

Execution of a Selection Operation

- **Selection** ($\sigma_{\text{conditions}} R$)
 - Assuming to select **one** tuple satisfying the specified conditions
- If there is **no index**, nor the target attribute is not a key
 - Comparisons: $\{R\}$
 - Disk I/Os: $|R|$



2019/7/11

Advance Data Engineering (©H.Yokota)

124

Pseudo Code for Selection

```

for (i = 1; i < |R|; i++) {
  read i-th page from disks into a buffer;
  f := true;
  while (f) {
    get a tuple t from the buffer;
    if t is empty then f := false;
    else {
      get the target attribute value v from t;
      if v satisfies the condition then output t;
    }
  }
}

```

Annotations in the original image:

- A blue box labeled "Read |R| pages in total" points to the "read i-th page from disks into a buffer;" line.
- A blue box labeled "Compare {R} tuples in total" points to the "if v satisfies the condition then output t;" line.

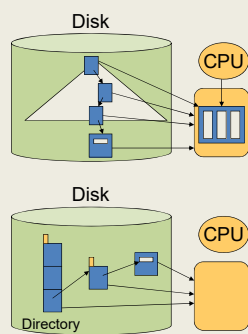
2019/7/11

Advance Data Engineering (©H.Yokota)

125

Selection with Indices

- If there is a index for the target attribute
- **B⁺-tree**
 - Disk I/Os: Height of the Tree + 1
 - Comparisons: Height of the Tree (H) x Search in a node by binary search ($\log_2 F^*$)
- **Extensible Hashing**
 - Disk I/Os: Directory + Local hash table + 1
 - Comparisons: constant



2019/7/11

Advance Data Engineering (©H.Yokota)

126

Range Queries

- In many cases, we need tuples satisfying conditions expressed as a range, such as $10 < x < 20$
- B⁺-tree is suitable for the range queries
 - Links between leaves are used for traversing adjacent tuples
 - At first, search for a tuple much with the lower bound (10), then traverse links to search for another tuple much with the higher bound (20)
 - Tuples between these bounds are the results
- Hash based indices are not effective for the range queries

2019/7/11

Advance Data Engineering (©H.Yokota)

127

Projection Operations

- **Projection** ($\pi_{\text{attribute-list } R}$)
 - At first, we focus on the **Delta Projection** to concentrate on the cost for attribute extraction
 - We can ignore the cost for eliminating duplication
 - Then we consider the execution of **Duplication Elimination** for Projection
 - Three types of algorithms: Nested Loop, Sort, Hash
- **Delta Projection**
 - All pages should be read, and all tuples in those pages should be handled to extract the target attributes

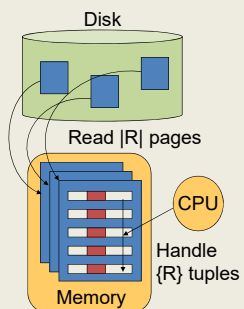
2019/7/11

Advance Data Engineering (©H.Yokota)

128

Delta Projection

- **Costs:**
 - Almost same as it for the selection except comparisons
 - Disk I/Os: $|R|$
 - Handled tuples: $\{R\}$
- Index has no effect for Projection operations
 - All tuples has to be accessed to extract desired attributes



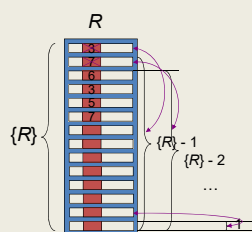
2019/7/11

Advance Data Engineering (©H.Yokota)

129

Nested Loop base Duplication Elimination (NLDE)

- At first, we concentrate on a straightforward algorithm ignoring disk I/Os
- Search the same value from the top by scanning rest tuples
- Comparisons:
 - $(\{R\} - 1) + (\{R\} - 2) + \dots + 1$
 - $(\{R\} \times (\{R\} - 1)) / 2$



2019/7/11

Advance Data Engineering (©H.Yokota)

130

Pseudo Code for NLDE

```

for (i = 1; i ≤ {R}; i++) {
  f := true;
  get i-th tuple  $t_x$  and attribute value  $v_x$ ;
  for (j = i + 1; j ≤ {R}; j++) {
    get j-th tuple  $t_y$  and attribute value  $v_y$  in  $t_y$ ;
    if  $v_x == v_y$  then  $f := false$ ;
  }
  if  $f$  then output  $t_x$ 
}

```

Outer Loop

Inner Loop

- This is a very straightforward algorithm and you can apply many optimizations on it

2019/7/11

Advance Data Engineering (©H.Yokota)

131

Costs for NLDE with Disk I/Os

- The costs depend on the size of buffer
- If the buffer size is enough to keep all disk pages into the memory:
 - Disk I/Os: $|R|$
- If the buffer size is just the same as a disk page:
 - We have to apply the nested loop algorithm for disk I/Os
 - Disk I/Os: $(|R| \times (|R| + 1)) / 2$

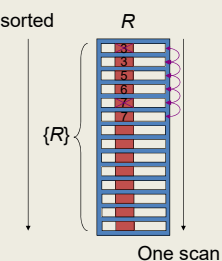
2019/7/11

Advance Data Engineering (©H.Yokota)

132

Sort based Duplication Elimination (SDE)

- If the tuples are sorted by the target attribute, it becomes very easy to check duplications
- Here, we first ignore disk I/Os again
- Compare continuous two tuples
- Duplications can be eliminated by one scan



2019/7/11

Advance Data Engineering (©H. Yokota)

133

Pseudo Code for SDE

```

for (i = 1; i ≤ {R}; i++) {
    get i-th tuple  $t_x$  and attribute value  $v_x$  in  $t_x$ ;
    if ( i < {R} ) then {
        get (i+1)-th tuple  $t_y$  and attribute value  $v_y$  in  $t_y$ ;
        if  $v_x \neq v_y$  then output  $t_x$ ;
    }
    else output  $t_x$ ;
}

```

One scan

- Costs to eliminate duplications: $\{R\}$
- If the sorted tuples are stored into disk pages, the number of disk I/Os is $|R|$
- But, it does not contain costs for sorting tuples

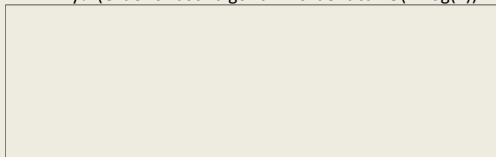
2019/7/11

Advance Data Engineering (©H. Yokota)

134

Sort Algorithms for DB Operations

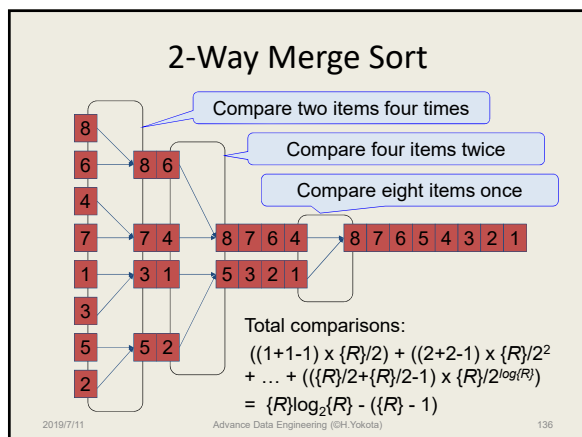
- There are many sort algorithms
 - Bubble Sort, Bitonic Sort, Quick Sort, N-Way Merge Sort, and so on
- The Quick Sort is not suitable for database operations containing disk I/Os, while the N-Way Merge Sort is commonly used for them
 - Why? (Order of both algorithm is identical: $O(n \log n)$)

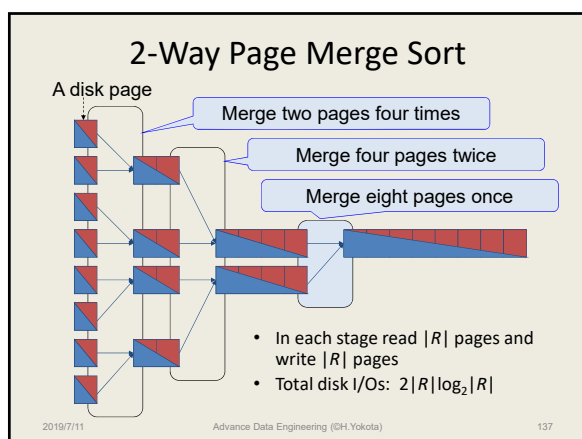


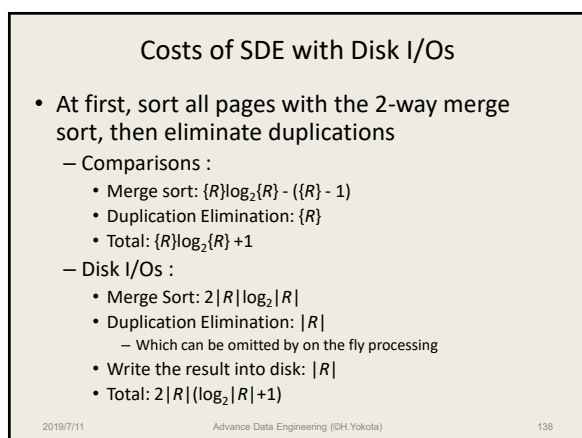
2019/7/11

Advance Data Engineering (©H. Yokota)

135







Hash based Duplication Elimination (HDE)

- Apply a hash function for each tuple
- If there is no hash collision, results can be derived by one scan (optimistic)
 - Comparison : $\{R\}$
 - Disk I/O : $|R|$
- When the size of hash table is small, many collisions occur
- When the size of hash table is large, it cannot be stored on memory
- The costs for generating large extensible hashing is high

2019/7/11 Advance Data Engineering (©H.Yokota) 139

Pseudo Code of Hash base DE

```

for (i = 1; i ≤ {R}; i++) {
  get i-th tuple t and attribute value v in t;
  x = h(v); Applying the hash function
  if ht[x] is empty then {
    ht[x] := t; Insert tuple in the hash table
    output t;
  }
}

```

* Strictly, we have to consider the treatment of Hash collisions

2019/7/11 Advance Data Engineering (©H.Yokota) 140

Costs for a Projection Operation

- Costs for duplication elimination (DE)
+ costs for extracting attributes (EA)

	Disk I/Os	Comparisons
NLDE + EA (large buffer)	$2 R $	$(\{R\} \times (\{R\} - 1)) / 2$
NLDE + EA (a page size buffer)	$(R \times (R + 1)) / 2 + R $	$(\{R\} \times (\{R\} - 1)) / 2$
SDE + EA	$ R (2\log_2 R +3)$	$\{R\}\log_2\{R\} + \{R\} - 1$
HDE + EA (Optimistic)	$2 R $	$\{R\}$

2019/7/11 Advance Data Engineering (©H.Yokota) 141

Join Operation

- Join is very time consuming operation
- Here we consider eq-join
- Eq-join is one of the most frequently used relational operation
 - Especially for the third normal form
- There are three algorithms
 - Nested Loop Join
 - Sort Merge Join
 - Hash Join
 - Hash Join cannot be used for θ -Join

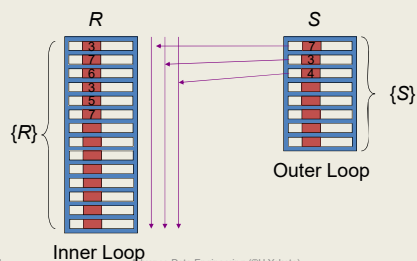
2019/7/11

Advance Data Engineering (©H.Yokota)

142

Nested Loop Join

- The most naïve algorithm
- Check all tuples in one relation (R) for each tuple in another relation (S)



2019/7/11

Advance Data Engineering (©H.Yokota)

143

Pseudo Code for Nested Join

```

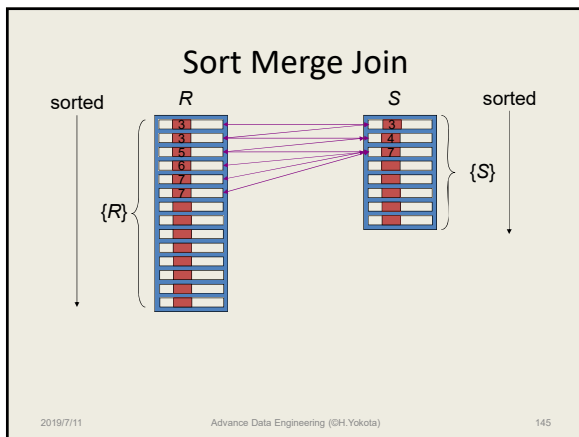
for (i = 1; i ≤ |S|; i++) {
  f := true;
  get i-th tuple tS and attribute value vS in tS;
  for (j = 1; j ≤ |R|; j++) {
    get j-th tuple tR and attribute value vR in tR;
    if vR == vS then output (tR, tS);
  }
}
  
```

- Comparisons : $|R| \times |S|$
- Disk I/Os : $|R| \times |S|$

2019/7/11

Advance Data Engineering (©H.Yokota)

144



Pseudo Code for Sort Merge Join

```

i := 1;
j := 1;
while (i ≤ {S} and j ≤ {R}) {
  get i-th tuple  $t_S$  and attribute value  $v_S$  in  $t_S$ ;
  get j-th tuple  $t_R$  and attribute value  $v_R$  in  $t_R$ ;
  if  $v_S == v_R$  then {
    output ( $t_R, t_S$ );
    if  $j ≤ \{R\}$  then  $j := j + 1$  else  $i := i + 1$ ;
  }
  else if ( $v_S < v_R$ ) then  $i := i + 1$ ;
  else  $j := j + 1$ ;
}

```

2019/7/11 Advance Data Engineering (©H.Yokota) 146

Costs for Sort Merge Join

- Comparisons (worst) : $\{R\} + \{S\} - 1$
- Costs for sorting are ignored
- Applying 2-way merge sort in advance
 - Comparisons : $\{R\}\log\{R\} + \{S\}\log\{S\} + 1$
- Disk I/Os:
 - 2-way merge sort: $2 \lceil R \rceil (\log_2 \lceil R \rceil + 1)$
 - Sort merge join: $\lceil R \rceil + \lceil S \rceil$
 - Disk I/O : $\lceil R \rceil (2\log \lceil R \rceil + 3) + \lceil S \rceil (2\log \lceil S \rceil + 3)$

2019/7/11 Advance Data Engineering (©H.Yokota) 147

