CS525: Advanced Database Organization

Notes 6: Query Processing Logical Optimization

Yousef M. Elmehdwi

Department of Computer Science

Illinois Institute of Technology

yelmehdwi@iit.edu

October 11, 2018

Slides: adapted from a courses taught by Hector Garcia-Molina, Stanford, Shun Yan Cheung, Emory University, & Jennifer Welch, Texas A&M, Ramon Lawrence & Introduction to Database Systems by ITL Education Solutions Limited

Basic Steps in Processing an SQL Query



- SQL \Rightarrow parse tree \Rightarrow expression of relational algebra (initial logical query plan)
- Today: consider ways of transformations to improve the query plan
 - Algebraic laws for improving query plans

- The translation rules converting a parse tree to a logical query tree do not always produce the best logical query tree.
- It is often possible to optimize the logical query tree by applying relational algebra laws to convert the original tree into a more efficient logical query tree.
- Next we'll survey some of these laws
- Optimizing a logical query tree using relational algebra laws is called heuristic optimization



- What are transformation rules?
 - preserve equivalence
- What are good transformations?
 - reduce query execution costs

- Two queries q₁ and q₂ are equivalent:
 - If for every database instance I (contents of all the tables)
 - Both queries have the same result
- $q_1 \equiv q_2$ iff \forall I: $q_1(I) = q_2(I)$

StarsIn(title. year, startName) MovieStar(name, address, gender, birthdate)



- Join (\bowtie) is commutative: $R \bowtie S = S \bowtie R$
- Join (\bowtie) is associative: (R \bowtie S) \bowtie T = R \bowtie (S \bowtie T)

- Carry attribute names in results, so order is not important
- Can also write as trees, e.g.:



- Different ordering in the execution of the ⋈ operation can produce different intermediate results (often with large difference in size of result sets)
- So one of the topics (problems) in query optimization will be:
 - \bullet Find the optimal join ordering of a set of \bowtie operations

- Cross product (\times) is commutative: R \times S = S \times R
- \bullet Cross product (×) is associative: (R \times S) \times T=R \times (S \times T)

- Union (U) is commutative: R U S = S U R
- Union (U) is associative: (R \cup S) \cup T = R \cup (S \cup T)

- Selections usually reduce the size of the relation (decrease the number of rows)
- Usually good to do selections early, i.e., "push them down the tree"
 - Perform selection as early as possible (but take existing indexes on relations into account)
- Also can be helpful to break up a complex selection into parts

• Selection is idempoten. (multiple applications of the same selection have no additional effect beyond the first one)

• $\sigma_p(\mathbf{R}) = \sigma_p \sigma_p(\mathbf{R})$

• Select operations are commutative. (the order selections are applied in has no effect on the eventual result)

•
$$\sigma_p \sigma_q(\mathbf{R}) = \sigma_q \sigma_p(\mathbf{R})$$

 The selection condition involving conjunction of two or more predicates can be deconstructed into a sequence of individual select operations.

•
$$\sigma_{p_1 \wedge p_2}(\mathbf{R}) = \sigma_{p_1}(\sigma_{p_2}(\mathbf{R})) = \sigma_{p_2}(\sigma_{p_1}(\mathbf{R}))$$

• This transformation is called cascading of select operator.

- R = {a,a,b,b,c}
- S = {b,b,c,c,d}
- $R \cup S = ?$
- Option 1: SUM
 - $R \cup S = \{a,a,b,b,b,b,c,c,c,d\}$
- Option 2: MAX
 - $R \cup S = \{a,a,b,b,c,c,d\}$

- Use ''SUM'' option for bag unions
- CAREFUL!. Some rules cannot be used for bags

Laws for Bags and Sets Can Differ

- Example of an Algebraic Law that holds for set, but not for bags
- We know from Set Theory that

A \cap_{set} (B \cup_{set} C) = (A \cap_{set} B) \cup_{set} (A \cap_{set} C)

- But, this law does not hold for bags:
 - Suppose bags A, B, and C were each $\{x\}$

$$A \cap_{bag} (B \cup_{bag} C) = \{x\} \cap_{bag} (\{x\} \cup_{bag} \{x\})$$
$$= \{x\} \cap_{bag} \{x, x\}$$
$$= \{x\}$$

$$(A \cap_{bag} B) \cup_{bag} (A \cap_{bag} C) = (\{x\} \cap_{bag} \{x\}) \cup_{bag} (\{x\} \cap_{bag} \{x\})$$
$$= \{x\} \cup_{bag} \{x\}$$
$$= \{x, x\}$$

- Push selections through the binary operators: product, union, intersection, difference, and join.
 - 1. Must push selection to both arguments:

• $\sigma_p(\mathbf{R} \cup \mathbf{S}) = \sigma_p(\mathbf{R}) \cup \sigma_p(\mathbf{S})$

2. Must push to first argument, optional for second:

$$\sigma_p(\mathbf{R} - \mathbf{S}) = \sigma_p(\mathbf{R}) - \sigma_p(\mathbf{S})$$

•
$$\sigma_p(R - S) = \sigma_p(R) - S$$

- 3. Push to at least one argument with all attributes mentioned in p:
 - product, natural join, theta join, intersection
 - e.g., $\sigma_p(\mathbf{R} \times \mathbf{S}) = \sigma_p(\mathbf{R}) \times \mathbf{S}$, if p contains only attributes from R

• If the condition p in $\sigma_p(R \cap S)$ is compound (p = p₁ and p₂), to split p up, we can use:

•
$$\sigma_{p_1 \wedge p_2}(\mathbf{R}) = \sigma_{p_1}(\sigma_{p_2}(\mathbf{R})) = \sigma_{p_2}(\sigma_{p_1}(\mathbf{R}))$$

Example

• R(a,b)

$$\sigma_{a=3\wedge c=4}(R\cap S) = \sigma_{a=3}(\sigma_{c=4}(R\cap S))$$
$$= \sigma_{a=3}(R\cap \sigma_{c=4}(S))$$
$$= \sigma_{a=3}(R) \cap \sigma_{c=4}(S)$$

Rules σ , \bowtie combined

• If the selection condition p involves only the attributes of R and q involves the attributes of S, then the select operation distributes.

•
$$\sigma_{p \wedge q}(\mathbf{R} \bowtie \mathbf{S}) = \sigma_p(\mathbf{R}) \bowtie \sigma_q(\mathbf{S})$$

Let

- p = predicate with only R attributes
- $\bullet \ q = predicate \ with \ only \ S \ attributes$
- m = predicate with R,S attributes
- $\sigma_p(\mathbf{R} \bowtie \mathbf{S}) = \sigma_p(\mathbf{R}) \bowtie \mathbf{S}$

•
$$\sigma_q(\mathbf{R} \bowtie \mathbf{S}) = \mathbf{R} \bowtie \sigma_p(\mathbf{S})$$

Some Rules can be Derived:

•
$$\sigma_{p \wedge q}(\mathbf{R} \bowtie \mathbf{S}) = \sigma_p(\mathbf{R}) \bowtie \sigma_q(\mathbf{S})$$

• $\sigma_{p \wedge q}(\mathbf{R} \bowtie \mathbf{S}) = \sigma_p(\sigma_q(\mathbf{R}) \bowtie \sigma_q(\mathbf{S}))$

• $\sigma_{p \wedge q \wedge m}(\mathbf{R} \boxtimes \mathbf{S}) = \sigma_m(\sigma_p(\mathbf{R}) \boxtimes \sigma_q(\mathbf{S}))$

Derivation for first one

$$\sigma_{p \wedge q}(R \bowtie S) = \sigma_p(\sigma_q(R \bowtie S))$$
$$= \sigma_p(R \bowtie \sigma_q(S))$$
$$= \sigma_p(R) \bowtie \sigma_q(S)$$

Pushing Selections

Example

- Employee(fname, salary, dno)
- Dept(dname, dno)
- $\sigma_{dname='Research'}(Employee \bowtie Dept) = Employee \bowtie \sigma_{dname='Research'}(Dept)$
- 'Pushing down'' a selection (σ) will result in a smaller intermediate result set



Example

- Employee(fname, salary, dno)
- Dept(dname, dno)
- $\sigma_{dname='Research \land fname='John'}$ (Employee \bowtie Dept)
 - $= \sigma_{\textit{fname}='\textit{John'}} (\sigma_{\textit{dname}='\textit{Research'}} (\texttt{Employee} \bowtie \texttt{Dept}))$
 - $= \sigma_{\textit{fname}='\textit{John'}}(\texttt{Employee} \bowtie \sigma_{\textit{dname}='\textit{Research'}}(\texttt{Dept}))$
 - $= \sigma_{fname='John'}$ (Employee) $\bowtie \sigma_{dname='Research'}$ (Dept)

- Simple query optimization
 - The running time of database operations depends on:
 - The size of the input relations (operands)
 - Therefore: It is always beneficial (for running time) to reduce the size of the input relation(s)

Reducing the size of input relation using σ

• The selection operator σ can reduce the size of the input relation of some operators



• The input relation of \bowtie in the second case $\sigma_{birthday}$ LIKE '%1960' (StarsIn) can be much smaller than the input relation StarsIn

Simple query optimization technique: "push select down"

- One of the many query optimization techniques used by the DBMS is execute a σ_p as soon as possible.
- In terms of a query tree, it means that the σ_p operation is push as far down the logical query tree as possible



Note: "push select down" query optimization technique

 When a query contains a virtual table, then the σ_p operation is pushed down the logical query tree as far as possible is not sufficient

Example

- Relations:
 - StarsIn(title, year, starName, birthday) // Movie stars
 - **Movies**(title, year, genre, studioName) // Movies

```
• View:
```

```
CREATE VIEW
MoviesOf1996 AS {
SELECT *
FROM Movies
WHERE year = 1996
```

• Corresponding logical query plan

$$\sigma_{year=1996}$$

Example (Continue)

- Query: Find all movie stars and their studio name in movies of 1996
 - SELECTstarName, studioNameFROMMoviesOf1996, StarsInWHEREMoviesOf1996.title = StarsIn.title
- initial logical query plan



Note: "push select down" query optimization technique

Example (Continue)

• After replacing the virtual table with the corresponding query:



• However, the optimal query plan is as follows:



Amendment to the simple query optimization technique

- If there are virtual table in the query plan, then to find the optimal query plan, we must
 - Push any selection σ operators in the virtual table as far up the query tree as possible
 - Push every selection σ operators in the resulting query tree as far down the query tree as possible

Example

• Query plan after incorporating the virtual table query:



Example (Continue)

• Use this algebraic law in the reverse order: $\sigma_p(\mathbb{R} \bowtie \mathbb{S}) = \sigma_p(\mathbb{R}) \bowtie \mathbb{S}$ to push the $\sigma_{vear=1996}$ operation up the tree



Example (Continue)

- Both relations have the attribute year
- Use this algebraic law in the forward order:
 - $\sigma_p(R \bowtie S) = \sigma_p(R) \bowtie \sigma_p(S)$ to **push** the $\sigma_{year=1996}$ operation **down** the tree



Laws Involving Projections: Use of π in query optimization

- The projection operation π can remove unnecessary attributes from intermediate results
- Common use the project operation π in query optimization:
 - The projection operator π can be added anywhere in the relational algebra expression (= logical query plan/tree), as long as:
 - $\bullet \ \pi$ will only eliminate attributes that are not used by an operator that is located high up the tree



- Consider adding in additional projections
- Adding a projection lower in the tree can improve performance, since often tuple size is reduced
 - Usually not as helpful as pushing selections down
- If a projection is inserted in the tree, then none of the eliminated attributes can appear above this point in the tree

- If a query contains a sequence of project operations, only the final operation is needed, the others can be omitted.
 - $\pi_{L_1}\Big(\pi_{L_2}\big(\ldots\big(\pi_{L_n}(R)\big)\ldots\big)\Big) = \pi_{L_1}(R)$, where $L_i \subseteq L_{i+1}$ for $i \in [1,n)$
 - This transformation is called cascading of project operator.

Let:

- X = set of attributes
- Y = set of attributes
- $XY = X \cup Y$
- $\pi_{XY}(R) = \pi_X(\pi_Y(R))$ is this correct?
Let:

- X = set of attributes
- Y = set of attributes

•
$$XY = X \cup Y$$

• $\pi_{XY}(R) = \frac{\pi_X(\pi_Y(R))}{\pi_X(R)}$

- It is also possible to push a projection below a selection.
- If the selection condition *p* involves only the attributes *a*₁, *a*₂,..., *a_n* that are present in the projection list, the two operations can be commuted.

•
$$\pi_{a_1,a_2,\ldots,a_n}(\sigma_p(\mathbf{R})) = \sigma_p(\pi_{a_1,a_2,\ldots,a_n}(\mathbf{R}))$$

• Rule: $\pi_L(\sigma_p(\mathbf{R})) = \pi_L(\sigma_p(\pi_M(\mathbf{R})))$, where M is all attributes used by L or p

- $\mathbf{x} = \mathsf{subset} \mathsf{ of } R \mathsf{ attributes}$
- z = attributes in predicate p (subset of R attributes)

•
$$\pi_x(\sigma_p(\mathbf{R})) = \sigma_p(\pi_x(\mathbf{R}))$$

- $\mathbf{x} = \mathsf{subset} \mathsf{ of } R \mathsf{ attributes}$
- z = attributes in predicate p (subset of R attributes)

•
$$\pi_x(\sigma_p(\mathbf{R})) = \sigma_p(\pi_x(\mathbf{R}))$$

- x = subset of R attributes
- z = attributes in predicate p (subset of R attributes)

•
$$\pi_{x}(\sigma_{p}(\mathbf{R})) = \pi_{x}(\sigma_{p}(\pi_{xz}(\mathbf{R})))$$

- x = subset of R attributes
- y = subset of S attributes
- z = intersection of R,S attributes
- $\pi_{xy}(\mathbf{R} \bowtie \mathbf{S}) = \pi_{xy}(\pi_{xz}(\mathbf{R}) \bowtie \pi_{yz}(\mathbf{S}))$

- x = subset of R attributes
- y = subset of S attributes
- z = intersection of R,S attributes

•
$$\pi_{xy}(\sigma_p(\mathbb{R} \boxtimes \mathbb{S})) = \pi_{xy}(\sigma_p(\pi_{xz'}(\mathbb{R}) \boxtimes \pi_{yz'}(\mathbb{S})))$$
, where $z' = z \cup \{attributes used in p\}$

Push Projection Below Selection?

Is it a good idea?

Example

- Relations:
 - StarsIn(title, movieYear, starName, birthday) // Movie stars
 - SELECT starName FROM StarsIn WHERE movieYear = 1996;



Extra work to scan through StarsIn twice

- Laws by definition: These are not really laws, but they are the definition of the ⋈ operator:
 - ($\mathbb{R} \bowtie_p \mathbb{S}$) = $\sigma_p(\mathbb{R} \times \mathbb{S})$ (theta join)
 - (R \bowtie S) = $\pi_L(\sigma_p(R \times S))$ (natural join)
 - where p equates same-name attributes in R and S, and L includes all attributes of R and S dropping duplicates
- To improve a logical query plan, replace a product followed by a selection with a join
 - Join algorithms are usually faster than doing product followed by selection on the (very large) result of the product

- Moving δ down the tree is potentially beneficial as it can reduce the size of intermediate relations
- Can be eliminated if argument has no duplicates
 - a relation with a primary key
 - a relation resulting from a grouping operator
- $\bullet\,$ Push the δ operation through product, join,theta-join, selection, and bag intersection
 - Ex: $\delta(\mathbf{R} \times \mathbf{S}) = \delta(\mathbf{R}) \times \delta(\mathbf{S})$
 - The result of δ is always a set (i.e.: no duplicates)
- $\bullet\,$ Cannot push $\delta\,$ through bag union, bag difference or projection
- The cost saving resulting from pushing down δ is usually small. Therefore, this optimization step is often not implemented

Example

- R has two copies of tuple t
- S has one copy of t
- T(a,b) contains only (1,2) and (1,3)
- Bag Union
 - $\delta(\mathbf{R} \cup_{\textit{bag}} \mathbf{S})$ has one copy of t
 - $\delta(\mathbf{R}) \cup_{\textit{bag}} \delta(\mathbf{S})$ has two copies of t

Bag difference

- δ (R S) has one copy of t
- $\delta(\mathbf{R}) \delta(\mathbf{S})$ has no copies of t

Bag projection

•
$$\delta(\pi_a(\mathbf{T})) = \{\mathbf{1}\}$$

• $\pi_a(\delta(T)) = \{1,1\}$

- The grouping operator only interact with very few relation algebra operations:
 - 1. γ_L produces a set, therefore the δ operation is unnecessary

• $\delta(\gamma_L(\mathbf{R})) = \gamma_L(\mathbf{R})$

2. You can project out some attributes as long as you keep the grouping attributes:

• $\gamma_L(\mathbf{R}) = \gamma_L(\pi_M(\mathbf{R}))$, where M must contain all attributes used by γ_L

- The aggregate functions max and min can tolerate removal of duplicates:
 - $\gamma_L(\mathbf{R}) = \gamma_L(\delta(\mathbf{R}))$, where $\gamma_L = \max$ or min
 - max(5,5,3)=max(5,3)

- Eliminate common sub-expressions
- Detect constant expressions

Applying the Algebraic laws for query optimization

Example

- MovieStar(name, addr, gender, birthdate)
- StarsIn(title, year, starName)
- Query: For each (movie) year, find the earliest birthday (youngest movie star) in that (movie) year

٩	SELECT	year, max(birthday)
	FROM	movieStar, StarsIn
	WHERE	name = starName
	GROUP BY	year

• The initial logical query plan is as follows:



Example (Continue)

• Apply: $\mathbb{R} \bowtie_p \mathbb{S} = \sigma_p(\mathbb{R} \times \mathbb{S})$, where p = "name = starName"



Example (Continue)

• Apply: $\gamma_L(\mathbf{R}) = \gamma_L(\delta(\mathbf{R}))$, where $\gamma_L = \max or \min$



Applying the Algebraic laws for query optimization

Example (Continue)

• Optionally, you can insert a projection at the top



Applying the Algebraic laws for query optimization

Example (Continue)

• Optionally, you can insert a couple of projections at the bottom



- Heuristic query optimization takes a logical query tree as input and constructs a more efficient logical query tree by applying equivalence preserving relational algebra laws.
- Equivalence preserving transformations insure that the query result is identical before and after the transformation is applied. Two logical query trees are equivalent if they produce the same result.
- Note that heuristic optimization does not always produce the most efficient logical query tree as the rules applied are only heuristics!

- 1. Deconstruct conjunctive selections into a sequence of single selection operations.
- 2. Move selection operations down the query tree for the earliest possible execution.
- 3. Replace Cartesian product operations that are followed by a selection condition by join operations.
- 4. Execute first selection and join operations that will produce the smallest relations.
- 5. Deconstruct and move as far down the tree as possible lists of projection attributes, creating new projections where needed

- No transformation is always good at the l.q.p level
- Selections
 - push down tree as far as possible
 - if condition is an AND, split and push separately
 - sometimes need to push up before pushing down
- Projections
 - can be pushed down
 - new ones can be added (but be careful)
- Duplicate elimination
 - sometimes can be removed
- Selection/product combinations
 - can sometimes be replaced with join
- Many transformations lead to "promising" plans

- Relational algebra level
 - transformations
 - good transformations
- Detailed query plan level
 - estimate costs
 - generate and compare plans

- A canonical logical query tree is a logical query tree where all associative and commutative operators with more than two operands are converted into multi-operand operators.
 - This makes it more convenient and obvious that the operands can be combined in any order.
- This is especially important for joins as the order of joins may make a significant difference in the performance of the query

- The transformations discussed so far intuitively seem like good ideas
- But how can we evaluate them more scientifically?
- Estimate size of relations, also helpful in evaluating physical query plans

- Logical query plan
 - a query tree where the nodes consist of relational algebra operators
- Physical query plan
 - a query tree where the nodes consist of relational algebra algorithms
 - There are different (implementation) algorithms for a relational algebra operator
 - Each with different cost (# disk IOs) and memory requirement

Query Optimization Physical Query Plan

• Physical query plan is derived from a logical query plan by:

- 1. Selecting an order and grouping for operations like joins, unions, and intersections.
- 2. Deciding on an algorithm for each operator in the logical query plan.
 - e.g. For joins: Nested-loop join, sort join or hash join
- 3. Adding additional operators to the logical query tree such as sorting and scanning that are not present in the logical plan.
- 4. Determining if any operators should have their inputs materialized for efficiency.
- Whether we perform cost-based or heuristic optimization, we eventually must arrive at a physical query tree that can be executed by the evaluator.

- To determine when one physical query plan is better than another, we must have an estimate of the cost of the plan.
- Heuristic optimization is normally used to pick the best logical query plan.
- Cost-based optimization is used to determine the best physical query plan given a logical query plan.
- Note that both can be used in the same query processor (and typically are). Heuristic optimization is used to pick the best logical plan which is then optimized by cost-based techniques



- 1. We start with an initial logical query plan (obtained by transforming the parse tree into a relational algebra tree)
- 2. We transform this initial logical query plan into optimal logical query plan using Algebraic Laws
- 3. We choose the best feasible algorithm for each relational operator in the optimal logical query plan to obtain the optimal physical query plan
 - we will learn to find optimal logical query plan

Comparing different logical query plans

- Before we can improve a query plan, we must have a measure to let us tell the difference (in cost) between the different logical query plans
- Measuring the cost of logical query plans
 - 1. The ultimate cost measure is Execution time (#disk IOs performed) of the query plan
 - However, Execution time is a measure used for implementation algorithms
 - I.e.: the physical query plan
 - We are comparing different logical query plan
 - A good approximation of the excution time (# disk IOs) measure is the size (# tuples) of the result produced by the operations

- The answer to the question is determined by:
 - The size (# tuples) of the intermediate result relations produced by each logical query plan
 - Because, the size (# tuples) will determine the number of disk IO performed by the relational operators (algorithms) further up in the query tree
- We need a method to compute (estimate) the size of the intermediate results of the relational operators on the logical query plan

- The size (# tuples) of the result set of a relational operator is not dependent on the implementation algorithm.
- The differences between the algorithms are
 - running time
 - memory requirement
- ullet \Rightarrow The size of the result in the intermediate outputs will
 - Depend only on the order of the operations in the logical query plan
 - Does not depend on algorithm used to compute the result
- Thus, # tuples in the intermediate result of the query plan is a good estimate for the cost of the logical query plan

Steps to find optimal (logical) query plan

 Use the relational algebra Laws to find least cost logical query plan without considering the ordering of the join operations (the query plan has a smaller # tuples in the intermediate results)



Steps to find optimal (logical) query plan

 If there are more than 2 input relations, then, find the ordering of the join operations that results in the smallest # tuples in the intermediate results in the join tree



Steps to find optimal (logical) query plan

• Notice that the end result of all the joins are equal

Example (Continue)



• The only difference is the intermediate result sets

- Estimates of cost are essential if the optimizer is to determine which of the many query plans is likely to execute fastest
 - Estimating size of results (Operation Cost)
 - $\bullet~\mbox{Estimating}~\#~\mbox{of IOs}$
- Note that the query optimizer will very rarely know the exact cost of a query plan because the only way to know is to execute the query itself!
 - Since the cost to execute a query is much greater than the cost to optimize a query, we cannot execute the query to determine its cost!
- It is important to be able to estimate the cost of a query plan without executing it based on statistics and general formulas.
Estimating result size (Operation Cost)

- Statistics/Information about relations and attributes
 - **T(R)** = number of tuples in the relation R
 - **S(R)** = size (# of bytes) of a tuple of R
 - **B(R)** = number of blocks used to hold all tuples of relation R
 - V(R,A) = number of distinct values for attribute A



- A: 20 bytes String
- B: 4 bytes integer
- C: 8 bytes date
- D: 5 bytes String

•
$$T(R) = 5$$

- S(R) = 37 bytes
- V(R,A) = 3, V(R,B) = 1
 V(R,C) = 5, V(R,D) = 4

- Calculating the size of a relation after the projection operation is easy because we can compute it directly
- Recall: π does not remove duplicate values
- This can be exactly computed
- Every tuple changes size by a known amount.
- Estimating $S = \pi_a(R)$
 - $T(\pi_a(R)) = T(R)$
 - Number of tuples is unchanged

Estimating the (size of the) result set of a projection (π)

Example

- R(A,B,C) is a relation with A and B integers of 4 bytes each; C a string of 100 bytes.
- Tuple headers are 12 bytes.
- Blocks are 1024 bytes and have headers of 24 bytes.
- T(R) = 10,000 and B(R) = 1250.
- How many blocks do we need to store $U = \pi_{A,B}(\mathbf{R})$?

Answer

- T(U) =T(R)=10,000
- S(U) =12+4+4=20 bytes
- We can hence store $\frac{(1024-24)}{20} = 50$ tuples in one block.
- \therefore B(U)= $\frac{T(U)}{50} = \frac{10,000}{50} = 200$ blocks
- This projection shrinks the relation by a factor slightly more than 6

•
$$T(R_1 \times R_2) = T(R_1) \times T(R_2)$$

• $S(R_1 \times R_2) = S(R_1) + S(R_2)$

- Generally reduce the number of tuples, although the sizes of tuples remain the same
- General formula
 - $T(\sigma_p(R)) = T(R) \times sel_p(R)$, where $sel_p(R)$ is the estimated fraction of tuples in R that satisfy predicate p
 - i.e., $sel_p(R)$ is the estimated probability that a tuple in R satisfies p.
- How we calculate sel_p(R) depends on what p is.

- $sel_{A=c}(\mathbf{R}) = \frac{1}{V(R,A)}$
- Intuition:
 - There are V(R,A) distinct A-values in R.
 - Assuming that A-values are uniformly distributed, the probability that a tuple has A-value c is $\frac{1}{V(R,A)}$
- \therefore T($\sigma_{A=c}(\mathbf{R})$) = $\frac{T(R)}{V(R,A)}$, i.e., original number of tuples divided by number of different values of A

- R(A,B,C) is a relation.
- T(R) = 10,000
- V(R, A) = 50
- Estimate $T(\sigma_{A=10}(R))$

•
$$T(\sigma_{A=10}(R)) = \frac{T(R)}{V(R,A)} = \frac{10,000}{50} = 200$$

- Assumption:
 - Values in select expression A=constant are uniformly distributed over possible V(R,A) values.
- Alternate Assumption:
 - Values in select expression A=constant are uniformly distributed over domain with DOM(R,A) values.

- Better selectivity estimates are possible if we have more detailed statistics
- A DBMS typically collects histograms that detail the distribution of values.
- A histogram can be of two types:
 - equi-width histogram: the range of values is divided into equal-sized subranges.
 - equi-depth histograms: the sub ranges are chosen in such a way that the number of tuples within each sub range is equal.
- Such histograms are only available for base relations, however, not for sub-results.

- If a histogram is available for the attribute A, the number of tuples can be estimated with more accuracy.
- The range in which the value c belongs is first located in the histogram.
- |B|: number of values per bucket (# distinct values appearing in that range)
- #B: number of records in bucket

•
$$T(\sigma_{A=c}(R)) = \frac{\#B}{|B|}$$

- R(A,B,C) is a relation.
- T(R) = 10,000
- V(R, A) = 50
- Estimate $T(\sigma_{A=10}(R))$
- The DBMS has collected the following equi-width histogram on A

range	[1,10] [11,20]	[21,30]	[31,40]	[41,50]
tuples in range	50 2000	2000	3000	2950
$T(\overline{\sigma_{A=10}(R)}) = \frac{\#B}{ B } =$	$\frac{50}{10} = 5$			

- $sel_{A < c}(\mathbb{R}) = \frac{1}{3}$
- Intuition:
 - On average, you would think that the value should be $\frac{T(R)}{2}$. However, queries with inequalities tend to return less than half the tuples, so the rule compensates for this fact.
 - i.e., Queries involving an inequality tend to retrieve a small fraction of the possible tuples (usually you ask about something that is true of less than half the tuples)

- R(A,B,C) is a relation.
- T(R) = 10,000
- $T(\sigma_{A<10}(R)) = T(R) \times \frac{1}{3} \approx 3334$

Result size estimation: Estimate values in range: $\sigma_{A < c}(R)$ with c a constant

Example

- R(A,B,C) is a relation.
- T(R) = 10,000
- The DBMS statistics show that the values of the A attribute lie within the range [8, 57], uniformly distributed.
- Question: what would be a reasonable estimate of sel_{A<10} (R)?

Answer

- We see that 57- 8+1 different values of A are possible
- however only records with values A=8 or A=9 satisfy the filter A<10.
- Therefore, $sel_{A<10}(R) = \frac{2}{(57-8+1)} = \frac{2}{50} = 0.04$
- And hence, $T(\sigma_{A<10}(R)) = T(R) \times sel_{A<10}(R) = 400$

- S= $\sigma_{A\neq c}(\mathbf{R})$
- Fact:
 - $\sigma_{A \neq c}(\mathbf{R}) \cup \sigma_{A=c}(\mathbf{R}) = \mathbf{R}$
 - $\Leftrightarrow \sigma_{A \neq c}(\mathbf{R}) = \mathbf{R} \sigma_{A = c}(\mathbf{R})$
- Therefore,

 - $sel_{A \neq c}(\mathbf{R}) = \frac{V(R,A)-1}{V(R,A)}$ $T(S) = T(\mathbf{R}) \times \frac{V(R,A)-1}{V(R,A)}$

Result size estimation: 4. $\sigma_{\neg p}(R)$

•
$$sel_{\neg p}(\mathbf{R}) = 1 - sel_p(\mathbf{R})$$

Result size estimation: 5. $\sigma_{P_1 \wedge P_2}(R)$

- Simple selection clauses can be connected using AND or OR.
- $sel_{P_1 \wedge P_2}(\mathbf{R}) = sel_{p_1}(\mathbf{R}) \times sel_{P_2}(\mathbf{R})$
- Assumption: The conditions P_1 and P_2 are (statistically) independent
- Treat $\sigma_{P_1 \wedge P_2}(R)$ as $\sigma_{P_1}(\sigma_{P_2}(R))$ (Cascade of simple selections)
- The order does not matter, treating this as $\sigma_{P_2}(\sigma_{P_1}(\mathbf{R}))$ gives the same results.

Example

- R(A,B,C) is a relation. T(R) = 10,000. V(R,A) = 50
- Estimate the size of the result set S = $\sigma_{A=10 \land B < 20}$ (R)

Answer

•
$$sel_{A=10}(R) = \frac{1}{5}$$

- $sel_{B<20}(R) = \frac{1}{3}$
- T(S) = $sel_{A=10} \times sel_{B<20} \times T(R) = \frac{1}{50} \times \frac{1}{3} \times 10,000 = 66.67$

Result size estimation: 6. $\sigma_{P_1 \vee P_2}(R)$

•
$$P_1 \vee P_2 = \neg (\neg P_1 \land \neg P_2)$$

- Treat $\sigma_{P_1 \lor P_2}(R)$ as $\sigma_{\neg(\neg P_1 \land \neg P_2)}(R)$
- $sel_{P_1 \vee P_2}(\mathbf{R}) = 1 (1 sel_{p_1}(\mathbf{R})) \times (1 sel_{P_2}(\mathbf{R}))$

Example

- R(A,B,C) is a relation. T(R) = 10,000. V(R,A) = 50
- Estimate the size of the result set S = $\sigma_{A=10\vee B<20}$ (R)

Answer

•
$$sel_{A=10}(R) = \frac{1}{50}$$

•
$$sel_{B<20}(R) = \frac{1}{3}$$

• T(S) =
$$\left(1 - \left(1 - \frac{1}{50}\right)\left(1 - \frac{1}{3}\right)\right) \times T(R)$$

Result size estimation: R \bowtie S

- We will only study estimating the size of natural join.
 - Other types of joins are equivalent or can be translated into a cross-product followed by a selection.
- Assume the relation schema R(X,Y) and S(Y,Z), we join on $Y(R(X,Y) \bowtie S(Y,Z))$.
- Question: Estimate the size of (R(X,Y) ⋈ S(Y,Z))
- The challenge is we do not know how the set of values of Y in R relate to the values of Y in S. There are some possibilities:
 - If the Y attribute values in R(X,Y) and S(Y,Z) are disjoint

• $T(R(X,Y) \bowtie S(Y,Z)) = 0$

• If Y attribute is a key in S and a foreign key of R, so each tuple of R joins with exactly one tuple of S

• $T(R(X,Y) \bowtie S(Y,Z)) = T(R)$

• If almost every tuple in R and S has the same Y attribute value

• $T(R(X,Y) \bowtie S(Y,Z)) = T(R) \times T(S)$

• Range of T(R \bowtie S): 0 \leq T(R \bowtie S) \leq T(R) \times T(S)

Result size estimation: $R \bowtie S$: Simplifying Assumptions

- Without any assumptions on the joining attribute values, it is not possible to provide an estimation on the result T(R ⋈ S)
- Assumptions that helps use find an estimate of $R(X,Y) \bowtie S(Y,Z)$
- 1. The containment of value sets assumption
 - An attribute Y in a relation R(...,Y) always takes on a prefix of a fixed list of values: y₁ y₂ y₃ y₄ ...

Example	l
Relations:	l
• R(, Y)	l
• S(, Y)	l
• U(, Y)	l
• Attr values of Y in R can be one of: $y_1 \ y_2 \ \dots \ y_R$	l
• Attr values of Y in S can be one of: $y_1 y_2 \dots y_s$	l
• Attr values of Y in U can be one of: $y_1 \ y_2 \ \dots \ y_U$	

• Containment of value sets assumption will help to estimate the size of T(R \bowtie S)

- Assumptions that helps use find an estimate of $R(X,Y) \bowtie S(Y,Z)$
- 2. The preservation of value sets assumption
 - The join operation R(X,Y) ⋈ S(Y,Z) will preserve all the possible values of the non-joining attributes
 - In other words
 - The attribute values taken on by X in R(X,Y) \Join S(Y,Z) and R(X,Y) are same
 - The attribute values taken on by Z in R(X,Y) \bowtie S(Y,Z) and S(Y,Z) are same
 - \bullet preservation of value sets assumption will help to estimate the size of T(R \bowtie S \bowtie U)

Result size estimation: $R \bowtie S$ when joining on 1 attribute

- We can estimate the size of $R(X,Y) \bowtie S(Y,Z)$ as follows:
- Case 1. $V(R,Y) \ge V(S,Y)$
 - The tuples in relations R and S take on the following attribute values for the Y attribute:
 - Attr values of Y in R: $y_1 y_2 \dots y_{V(R,Y)}$
 - Attr values of Y in S: $y_1 y_2 \dots y_{V(S,Y)}$
 - Then every tuple t of S has a chance $\frac{1}{V(R,Y)}$ of joining with a given tuple of R.
 - There are T(R) tuples in R, therefore, one tuple t \in S will produce $\frac{T(R)}{V(R,Y)}$ number of matches
 - There are T(S) tuples in S, then estimated size of R \bowtie S is $\frac{T(R) \times T(S)}{V(R,Y)}$
- Case 2. $V(S,Y) \ge V(R,Y)$
 - estimated size of $\mathbb{R} \bowtie \mathbb{S}$ is $\frac{T(R) \times T(S)}{V(S,Y)}$

Result size estimation: R $\,\bowtie\,$ S when joining on 1 attribute

- In general, we divide by whichever of V(R,Y) and V(S,Y) is larger. That is:
- T(R \bowtie S) = $\frac{T(R) \times T(S)}{\max(V(R,Y),V(S,Y))}$

Example				
	R(a,b)	S(b,c)	U(c,d)	
	T(R)=1000 V(R,b)=20	T(S)=2000 V(S,b)=50	T(U)=5000	
		V(S,c)=100	V(U,c)=500	
• Estimate the size of $R \bowtie S \bowtie U$?				

Method 1: (ordering 1)

• R(a,b) \bowtie S(b,c) \bowtie U(c,d) = (R(a,b) \bowtie S(b,c)) \bowtie U(c,d)

• $T(R(a,b) \bowtie S(b,c)) = \frac{T(R) \times T(S)}{\max(V(R,b),V(S,b))} = \frac{1000 \times 2000}{\max\{20,50\}} = 40,000$

• The estimate of the size the join $(R(a,b) \bowtie S(b,c)) \bowtie U(c,d)$ is $= \frac{T(R(a,b) \bowtie S(b,c)) \times T(U)}{\max \left(V(R(a,b) \bowtie S(b,c),c), V(U,c) \right)}$

• From the preservation of value sets assumption, we have: $V(R(a,b) \bowtie S(b,c),c) = V(S,c)$, where V(S,c)=100 according to data

• :: T(R
$$\bowtie$$
S \bowtie U) = $\frac{T(R(a,b) \bowtie S(b,c)) \times T(U)}{\max \left(V(R(a,b) \bowtie S(b,c),c), V(U,c) \right)} = \frac{40,000 \times 5,000}{\max(100,500)} = 400,000$

Method 2: (ordering 2)

• $R(a,b) \bowtie S(b,c) \bowtie U(c,d) = R(a,b) \bowtie (S(b,c) \bowtie U(c,d))$

- $T(S(b,c) \bowtie U(c,d)) = \frac{T(S) \times T(U)}{\max(V(S,c),V(U,c))} = \frac{2000 \times 5000}{\max\{100,500\}} = 20,000$
- The estimate of the size the join R(a,b) \bowtie (S(b,c) \bowtie U(c,d)) is $=\frac{T(R) \times T(S(b,c) \bowtie U(c,d))}{\max\left(V(R,b), V(S(b,c) \bowtie U(c,d),b)\right)}$
- From the preservation of value sets assumption, we have: $V(S(b,c) \bowtie U(c,d), b) = V(S,b)$, where V(S,b)=50 according to data

• :: T(R
$$\bowtie$$
S \bowtie U) = $\frac{T(R) \times T(S(b,c) \bowtie U(c,d))}{\max(V(R,b), V(S(b,c) \bowtie U(c,d),b))} = \frac{1,000 \times 20,000}{\max(20,50)} = 400,000$

- Assume the relation schema R(X,Y₁,Y₂) and S(Y₁,Y₂,Z), i.e., we join on Y₁ and Y₂.
- General formula:

 $T(R(X,Y_1,Y_2) \bowtie S(Y_1,Y_2,Z)) = \frac{T(R) \times T(S)}{\max(V(R,Y_1),V(S,Y_1))\max(V(R,Y_2),V(S,Y_2))}$

Example			
	R(a,b)	S(b,c)	U(c,d)
	T(R)=1000	T(S)=2000	T(U)=5000
	V(R,b)=20	V(S,b)=50	
		V(S,c)=100	V(U,c)=500

• Estimate the size of $R \bowtie S \bowtie U$?

• Computed using this ordering: $R(a,b) \bowtie S(b,c) \bowtie U(c,d) = (R(a,b) \bowtie U(c,d)) \bowtie S(b,c)$

- A join operation with no common attributes will degenerates into a Cartesian product
- Example: $R(a,b) \bowtie U(c,d) \Rightarrow R(a,b) \times U(c,d)$

Method 3: (ordering 3)

- R(a,b) \bowtie S(b,c) \bowtie U(c,d) = (R(a,b) \bowtie U(c,d)) \bowtie S(b,c)
- $T(R(a,b) \bowtie U(c,d)) = T(R(a,b) \times U(c,d)) = 1000 \times 5000 = 5,000,000$
- The estimate of the size the join $(R(a,b) \bowtie U(c,d)) \bowtie S(b,c)$ is $= \frac{T(R(a,b) \bowtie U(c,d)) \times T(S)}{T(S)}$

 $\overline{\max\left(V\left(R(a,b)\bowtie U(c,d),b\right),V(S,b)\right)\times\max\left(V\left(R(a,b)\bowtie U(c,d),c\right),V(S,c)\right)}$

 From the preservation of value sets assumption, we have: V(R(a, b) ⋈ U(c, d), b) = V(R,b), V(R,b)=20 according to data V(R(a, b) ⋈ U(c, d), c) = V(U,c), V(U,c)=500 according to data

 ∴ T(R⋈S⋈U) = 5,000,000×2,000 max(20,50)×max(500,100) =400,000

The 2 assumptions (containment and preservation of value sets) allows us to re-order the join-order without affecting the size of the result set estimation

Estimating Join Sizes: Estimating V(R, a)

- The database will keep statistics on the number of distinct values for each attribute a in each relation *R*, *V*(*R*, *a*).
- When a sequence of operations is applied, it is necessary to estimate *V*(*R*, *a*) on the intermediate relations.
- For our purposes, there will be three common cases:
 - a is the primary key of R then V(R, a) = T(R)
 ⇒ The number of distinct values is the same as the # tuples in R.
 - a is a foreign key of R to another relation S then V(R, a) = T(S)
 ⇒ In the worst case, the number of distinct values of a cannot be larger than the number of tuples of S since a is a foreign key to the primary key of S.
 - If a selection occurs on relation R before a join, then V(R, a) after the selection is the same as V(R, a) before selection.

 \Rightarrow This is often strange since V(R, a) may be greater than # of tuples in intermediate result! $V(R, a) \neq \#$ of tuples in result.

- Bag-based: T(R)+T(S)
- Set-based:
 - Range of the result set of R $\,\cup\,$ S:
 - $\max(T(R),T(S)) \leq T(R \cup S) \leq T(R) + T(S)$ $\max(T(R),T(S)): R \subseteq S \text{ or } S \subseteq R$ $T(R) + T(S): R \cap S = \emptyset$ • Recommended estimate for R $\cup S$ $T(R \cup S) = \max(T(R),T(S)) + \frac{1}{2} \times \min(T(R),T(S))$ i.e.: maximum + $\frac{1}{2} \times (\text{smaller size})$

 $\bullet\,$ Range of the result set of R $\,\cap\,$ S

•
$$0 \leq T(R \cap S) \leq min(T(R),T(S))$$

• 0:
$$\mathbf{R} \cap \mathbf{S} = \emptyset$$

•
$$\min(T(R),T(S))$$
: $R \subseteq S$ or $S \subseteq R$

 $\bullet\,$ Recommended estimate for R $\,\cap\,$ S

- $T(R \cap S) = \frac{1}{2} \times \min(T(R), T(S))$
- i.e.: the average of the min and max

• Range of the result set of R-S

•
$$\max(0,T(R)-T(S)) \leq T(R-S) \leq T(R)$$

• $\max(0,T(R)-T(S)): R \subseteq S \text{ or } S \subseteq R$

• T(R): R
$$\cap$$
 S = \emptyset

• Recommended estimate for R-S

- $T(R-S) = T(R) \frac{1}{2} \times T(S)$
- (Probabilistically speaking: 50-50 chance that a tuple in S is also in R)
- Note: if T(R) $\frac{1}{2}$ × T(S) \leq 0, then T(R-S) = 0 (estimate)

Result size estimation: $\delta(R,A)$

- Range of the result set of $\delta(\mathbf{R},\mathbf{A})$
 - 1 \leq T(δ (R,A)) \leq T(R)
 - 1: all tuples have same attribute value
 - T(R): all tuples have different attribute values
- Recommended estimate for $\delta(R,A)$
 - If the database maintains statistics on the attribute values: $T(\delta(R,A)) = V(R,A)$
 - If no statistics available, then we use this estimate: $T(\delta(R,A)) = \frac{1}{2} \times T(R)$
- Recommended estimate for $\delta(R,A,B)$
 - If the database maintains statistics on the attribute values: $T(\delta(R,A,B)) = V(R,A) \times V(R,B)$
 - If no statistics available, then we use this estimate:

 $T(\delta(R,A,B)) = \frac{1}{4} \times T(R)$

• Range of the result set of $\gamma_L(\mathbf{R})$

• 1
$$\leq$$
 T $(\gamma_L(R))$) \leq T (R)

- 1: all tuples have same attribute value
- T(R): all tuples have different attribute values for attribute L
- Recommended estimate for $T(\gamma_L(R))$
 - If the database maintains statistics on the attribute values: $T(\gamma_L(R)) = V(R,L)$
 - If no statistics available, then we use this estimate:

 $T(\gamma_L(R)) = \frac{1}{2}^{n(L)} \times T(R)$, where n(L) = number of attributes in the attribute list L

As should be clear by now, result size estimation is not an exact artDon't forget: Statistics must be kept up to date. (cost?)

- Size estimates can also be used during heuristic optimization.
- In this case, we are not deciding on a physical plan, but rather determining if a given logical transformation will make sense.
- By using statistics, we can estimate intermediate relation sizes (independent of the physical operator chosen), and thus determine if the logical transformation is useful.
- Estimating cost of query plan
 - Estimating size of results
 - Estimating # of IOs (next)
 - Operator Implementations
- Generate and compare plans