PURDUE UNIVERSITY | Department of Computer Science

CS 44800: Introduction To Relational Database Systems

Query Optimization

Prof. Chris Clifton

26 October 2021

Indiana Center for Database Systems ™

PURDUE UNIVERSITY
Department of Computer Science

Outline - Query Processing

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

Estimating result size

- Keep statistics for relation R
  - T(R) : # tuples in R
  - S(R) : # of bytes in each R tuple
  - B(R): # of blocks to hold all R tuples
  - V(R, A) : # distinct values in R for attribute A

Example Table Statistics

<u>Example</u>

R

A	B	C	D
cat	1	10	a
cat	1	20	b
dog	1	30	a
dog	1	40	c
bat	1	50	d

A: 20 byte string
B: 4 byte integer
C: 8 byte date
D: 5 byte string

T(R) = 5 S(R) = 37

V(R,A) = 3 V(R,C) = 5

V(R,B) = 1 V(R,D) = 4

PURDUE UNIVERSITY
Department of Computer Science

Size estimates for W = R1 x R2

- $T(W) = T(R1) \times T(R2)$
- $S(W) = S(R1) + S(R2)$

PURDUE UNIVERSITY
Department of Computer Science

Size estimate for W = σ

- $S(W) = S(R)$
- $T(W) = ?$

Example: Select

Example

R

A	B	C	D
cat	1	10	a
cat	1	20	b
dog	1	30	a
dog	1	40	c
bat	1	50	d

V(R,A)=3
V(R,B)=1
V(R,C)=5
V(R,D)=4

$$W = \sigma_{z=val}(R) \quad T(W) = \frac{T(R)}{V(R,Z)}$$

Assumption:

- Values in select expression Z = val are uniformly distributed over possible V(R,Z) values ?
- Values in select expression Z = val are uniformly distributed over domain with DOM(R,Z) values ?

PURDUE UNIVERSITY
Department of Computer Science

Example

R

A	B	C	D
cat	1	10	a
cat	1	20	b
dog	1	30	a
dog	1	40	c
bat	1	50	d

Alternate assumption
V(R,A)=3 DOM(R,A)=10
V(R,B)=1 DOM(R,B)=10
V(R,C)=5 DOM(R,C)=10
V(R,D)=4 DOM(R,D)=10

$W = \sigma_{z=val}(R)$ T(W) = ?

PURDUE UNIVERSITY
Department of Computer Science

C=val $\Rightarrow$ T(W) = (1/10)1 + (1/10)1 + ...
= (5/10) = 0.5

B=val $\Rightarrow$ T(W)= (1/10)5 + 0 + 0 = 0.5

A=val $\Rightarrow$ T(W)= (1/10)2 + (1/10)2 + (1/10)1
= 0.5

PURDUE UNIVERSITY

Department of Computer Science

Example

R

A	B	C	D
cat	1	10	a
cat	1	20	b
dog	1	30	a
dog	1	40	c
bat	1	50	d

Alternate assumption
V(R,A)=3 DOM(R,A)=10
V(R,B)=1 DOM(R,B)=10
V(R,C)=5 DOM(R,C)=10
V(R,D)=4 DOM(R,D)=10

$$W = \sigma_{z=val}(R) \qquad T(W) = \frac{T(R)}{DOM(R,Z)}$$

PURDUE UNIVERSITY

Department of Computer Science

Selection cardinality

SC(R,A) = average # records that satisfy equality condition on R.A

SC(R,A) = V(R,A) {

- _____ T(R)
- _______ T(R) DOM(R,A)

PURDUE UNIVERSITY
Department of Computer Science

What about $W = \sigma_{z \geq val}(R)$?

T(W) = ?

- Solution # 1:
  T(W) = T(R)/2

- Solution # 2:
  T(W) = T(R)/3

PURDUE UNIVERSITY
Department of Computer Science

Selection: Range Estimates

- Solution # 3: Estimate values in range

Example R

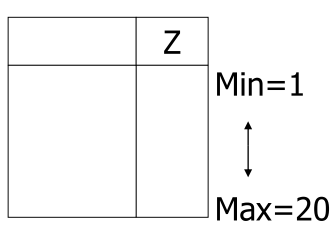

V(R,Z)=10

$W = \sigma_{z \geq 15}(R)$

$$f = \frac{20-15+1}{20-1+1} = \frac{6}{20}$$ (fraction of range)

$T(W) = f \times T(R)$

Equivalently:

$f \times V(R,Z)$ = fraction of distinct values

$$T(W) = [f \times V(Z,R)] \times T(R) = \frac{f \times T(R)}{V(Z,R)}$$

Estimating result size

- Keep statistics for relation R
  - T(R) : # tuples in R
  - S(R) : # of bytes in each R tuple
  - B(R): # of blocks to hold all R tuples
  - V(R, A) : # distinct values in R for attribute A

Size estimate for $W = R1 \bowtie R2$

- Let x = attributes of R1
- y = attributes of R2

Case 1

$X \cap Y = \emptyset$

S(W) same as R1 x R2

T(W)?

Case 2 $W = R1 \bowtie R2$ $X \cap Y = A$

R1

A	B	C

R2

A	D

Assumption:

$V(R1,A) \le V(R2,A) \Rightarrow$ Every A value in R1 is in R2

$V(R2,A) \le V(R1,A) \Rightarrow$ Every A value in R2 is in R1

"containment of value sets"

PURDUE UNIVERSITY
Department of Computer Science

Computing T(W) when $V(R1,A) \leq V(R2,A)$

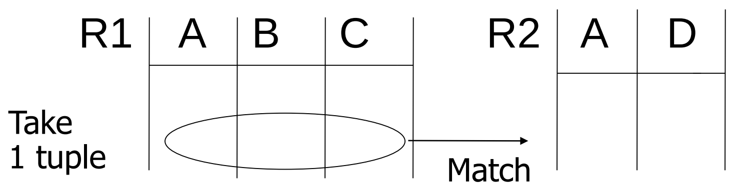

1 tuple matches with $\frac{T(R2)}{V(R2,A)}$ tuples...

so $$T(W) = \frac{T(R2)}{V(R2, A)} \times T(R1)$$

PURDUE UNIVERSITY
Department of Computer Science

$V(R1,A) \leq V(R2,A)$ $\quad T(W) = \frac{T(R2)\ T(R1)}{V(R2,A)}$

$V(R2,A) \leq V(R1,A)$ $\quad T(W) = \frac{T(R2)\ T(R1)}{V(R1,A)}$

[A is common attribute]

PURDUE UNIVERSITY
Department of Computer Science

In general W = R1 ⋈ R2

$$T(W) = \frac{T(R2)\ T(R1)}{\max\{ V(R1,A), V(R2,A) \}}$$

PURDUE UNIVERSITY
Department of Computer Science

Case 2 with alternate assumption

Values uniformly distributed over domain

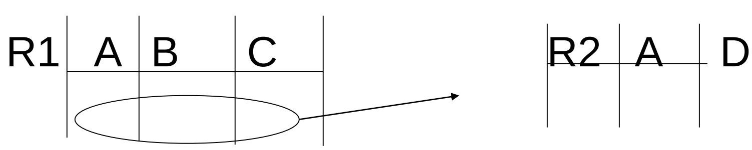

This tuple matches T(R2)/DOM(R2,A) so

$$T(W) = \frac{T(R2)\ T(R1)}{DOM(R2, A)} = \frac{T(R2)\ T(R1)}{DOM(R1, A)}$$

Assume the same

PURDUE UNIVERSITY
Department of Computer Science

Join: Tuple width

- In all cases:

  S(W) = S(R1) + S(R2) - S(A)

  size of attribute A

Size Estimation for Other Operations

- Projection: estimated size of $\prod_A(r)$ = $V(A,r)$
- Aggregation : estimated size of ${}_G\gamma_A(r)$ = $V(G,r)$
- Set operations
  - For unions/intersections of selections on the same relation: rewrite and use size estimate for selections
    - E.g., $\sigma_{\theta 1}(r) \cup \sigma_{\theta 2}(r)$ can be rewritten as $\sigma_{\theta 1 \text{ or } \theta 2}(r)$
  - For operations on different relations:
    - estimated size of $r \cup s$ = size of r + size of s.
    - estimated size of $r \cap s$ = minimum size of r and size of s.
    - estimated size of $r - s$ = r.
    - <u>All the three estimates may be quite inaccurate, but provide upper bounds on the sizes</u>.

Note: for complex expressions, need intermediate T,S,V results.

E.g. W = $[\sigma_{A=a}(R1)] \bowtie R2$

($\sigma_{A=a}(R1)$:) Treat as relation U

T(U) = T(R1)/V(R1,A)   S(U) = S(R1)

Also need V (U, *) !!

To estimate Vs

E.g., U = $\sigma_{A=a}(R1)$

Say R1 has attribs A,B,C,D

V(U, A) =

V(U, B) =

V(U, C) =

V(U, D) =

Example

R1

A	B	C	D
cat	1	10	10
cat	1	20	20
dog	1	30	10
dog	1	40	30
bat	1	50	10

V(R1,A)=3
V(R1,B)=1
V(R1,C)=5
V(R1,D)=3

$U = \sigma_{A=a}(R1)$

V(U,A) =1 V(U,B) =1 V(U,C) = $\frac{T(R1)}{V(R1,A)}$

V(D,U) ... somewhere in between

PURDUE UNIVERSITY
Department of Computer Science

Possible Guess $U = \sigma_{A=a}(R)$

V(U,A) = 1
V(U,B) = V(R,B)

For Joins $U = R1(A,B) \bowtie R2(A,C)$

V(U,A) = min { V(R1, A), V(R2, A) }

V(U,B) = V(R1, B)

V(U,C) = V(R2, C)

Example:

$Z = R1(A,B) \bowtie R2(B,C) \bowtie R3(C,D)$

R1	T(R1) = 1000	V(R1,A)=50	V(R1,B)=100
R2	T(R2) = 2000	V(R2,B)=200	V(R2,C)=300
R3	T(R3) = 3000	V(R3,C)=90	V(R3,D)=500

PURDUE UNIVERSITY
Department of Computer Science

Partial Result: $U = R \bowtie S$

$$T(U) = \frac{1000 \times 2000}{200}$$

V(U,A) = 50
V(U,B) = 100
V(U,C) = 300

PURDUE UNIVERSITY
Department of Computer Science

$Z = U \bowtie R3$

$$T(Z) = \frac{1000 \times 2000 \times 3000}{200 \times 300}$$

V(Z,A) = 50
V(Z,B) = 100
V(Z,C) = 90
V(Z,D) = 500

Estimation of Number of Distinct Values

Selections: $\sigma_\theta (r)$

- If θ forces A to take a specified value: $V(A,\sigma_\theta (r)) = 1$.
  - e.g., $A = 3$
- If θ forces A to take on one of a specified set of values:
  $V(A,\sigma_\theta (r))$ = number of specified values.
  - (e.g., $(A = 1 \vee A = 3 \vee A = 4)$),
- If the selection condition θ is of the form A *op* r
  estimated $V(A,\sigma_\theta (r)) = V(A.r) * s$
  - where s is the selectivity of the selection.
- In all the other cases: use approximate estimate of
  $\min(V(A,r), n_{\sigma\theta (r)})$
  - More accurate estimates can be made using probability theory, but this one works fine generally

Estimation of Distinct Values (Cont.)

Joins: $r \bowtie s$

- If all attributes in A are from r
  estimated $V(A, r \bowtie s) = \min(V(A,r), n_{r \bowtie s})$
- If A contains attributes $A1$ from r and $A2$ from s, then estimated
  $V(A,r \bowtie s)$ =
  $\min(V(A1,r)*V(A2 - A1,s), V(A1 - A2,r)*V(A2,s), n_{r \bowtie s})$
  - More accurate estimates can be made using probability theory, but this one works fine generally

Statistical Information for Cost Estimation

- n_r: number of tuples in a relation *r.*
- b_r: number of blocks containing tuples of *r.*
  - l_r: size of a tuple of *r.*
  - f_r: blocking factor of *r* — i.e., the number of tuples of *r* that fit into one block.
- *V(A, r):* number of distinct values that appear in *r* for attribute *A;* same as the size of $\prod_A(r)$.
- If tuples of *r* are stored together physically in a file, then:

$$b_r = \left\lceil \frac{n_r}{f_r} \right\rceil$$

More detailed V(A,r) statistics: Histograms

- Histogram on attribute *age* of relation *person*

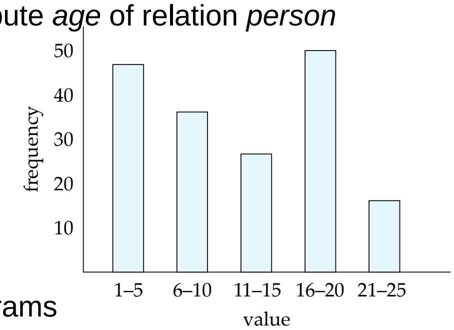

- **Equi-width** histograms
- **Equi-depth** histograms break up range such that each range has (approximately) the same number of tuples
  - E.g. (4, 8, 14, 19)
- Many databases also store *n* **most-frequent values** and their counts
  - Histogram is built on remaining values only