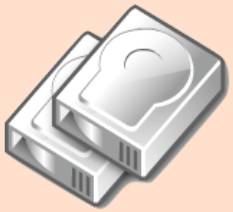
Cardinality Estimation

How Many Rows Does a Query Yield?





Cardinality Estimation

Database Profiles

Assumptions

Estimating Operator Cardinality

Selection σ

Projection π

Set Operations $\cup$, $\setminus$, $\times$

Join $\bowtie$

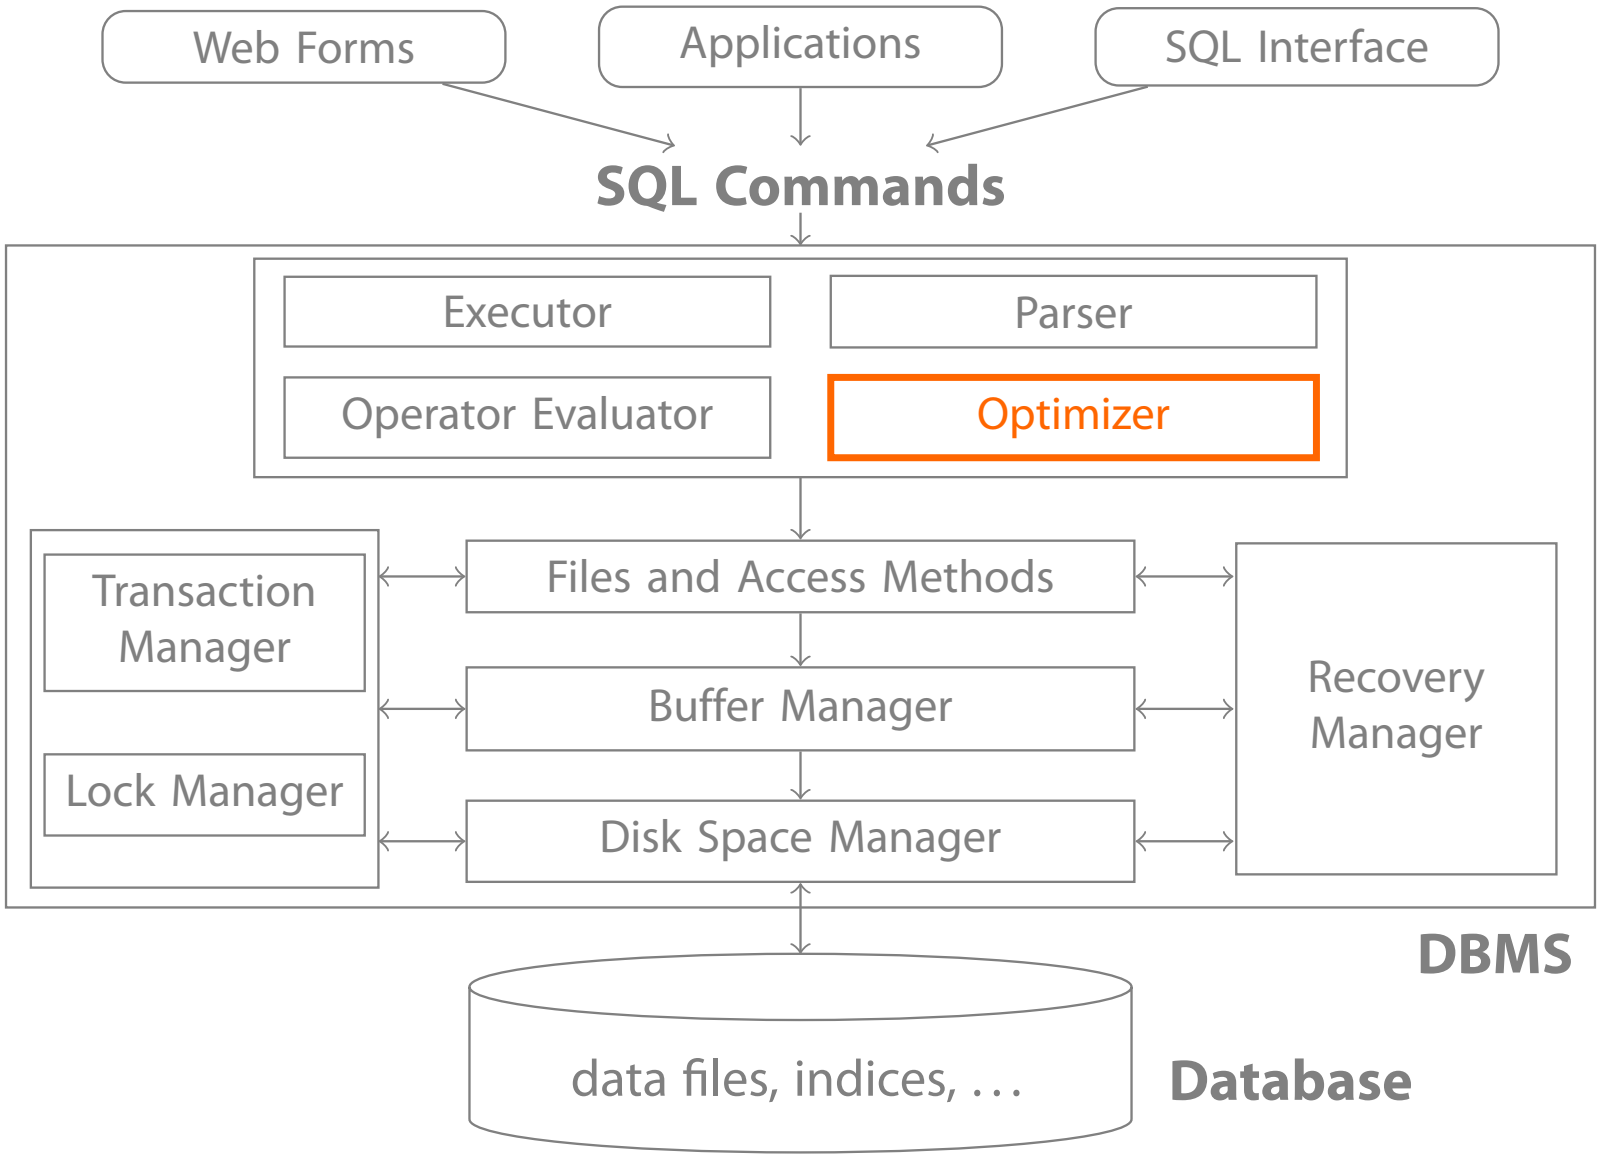
Histograms

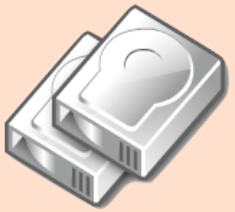
Equi-Width

Equi-Depth

Statistical Views

Cardinality Estimation





Cardinality Estimation

Database Profiles

Assumptions

Estimating Operator Cardinality

Selection σ

Projection π

Set Operations $\cup$, $\setminus$, $\times$

Join $\bowtie$

Histograms

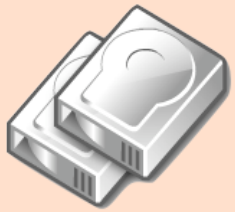
Equi-Width

Equi-Depth

Statistical Views

Cardinality Estimation

- A relational query optimizer performs a phase of **cost-based plan search** to identify the—presumably—"cheapest" alternative among a a set of equivalent execution plans (↗ Chapter on Query Optimization).
- Since page I/O cost dominates, the **estimated cardinality of a (sub-)query result** is crucial input to this search.
  - Cardinality typically measured in pages or rows.
- **Cardinality estimates** are also valuable when it comes to buffer "right-sizing" **before query evaluation starts** (*e.g.*, allocate B buffer pages and determine blocking factor b for external sort).



Cardinality Estimation
Database Profiles
Assumptions
Estimating Operator Cardinality
Selection σ
Projection π
Set Operations $\cup$, $\setminus$, $\times$
Join $\bowtie$
Histograms
Equi-Width
Equi-Depth
Statistical Views

Estimating Query Result Cardinality

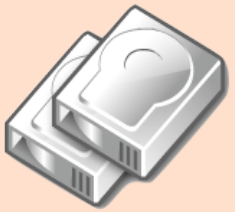
There are two principal approaches to query cardinality estimation:

1. **Database Profile.**
   Maintain statistical information about *numbers and sizes of tuples*, *distribution of attribute values* for **base relations**, as part of the **database catalog** (meta information) **during database updates**.

   - *Calculate* these parameters for **intermediate query results** based upon a (simple) statistical model during query optimization.

   - Typically, the statistical model is based upon the **uniformity** and **independence assumptions**.
   - Both are typically **not valid**, but they allow for simple calculations $\Rightarrow$ **limited accuracy**.
   - In order to improve accuracy, the system can record **histograms** to more closely model the actual value distributions in relations.





Cardinality Estimation
Database Profiles
Assumptions
Estimating Operator Cardinality
Selection σ
Projection π
Set Operations $\cup$, $\setminus$, $\times$
Join $\bowtie$
Histograms
Equi-Width
Equi-Depth
Statistical Views

Estimating Query Result Cardinality

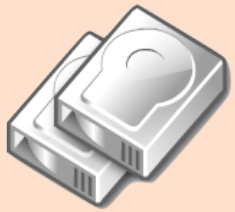
❷ Sampling Techniques.

Gather the necessary characteristics of a query plan (base relations and intermediate results) at **query execution time**:

- **Run query on a small sample** of the input.
- **Extrapolate** to the full input size.

- It is crucial to find the right balance between sample size and the resulting accuracy.

These slides focus on ❶ **Database Profiles.**



Cardinality Estimation

Database Profiles

Assumptions

Estimating Operator Cardinality

Selection σ

Projection π

Set Operations $\cup$, $\setminus$, $\times$

Join $\bowtie$

Histograms

Equi-Width

Equi-Depth

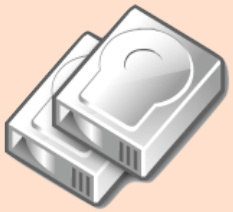
Statistical Views

Database Profiles

Keep profile information in the **database catalog**. Update whenever SQL DML commands are issued (database updates):

Typical database profile for relation R

$\lvert R \rvert$	**number of records** in relation R
N_R	**number of disk pages** allocated for these records
$s(R)$	average **record size**
$V(\mathtt{A}, R)$	**number of distinct values** of attribute $\mathtt{A}$
$\vdots$	*possibly many more*



Cardinality Estimation
Database Profiles
Assumptions
Estimating Operator Cardinality
Selection σ
Projection π
Set Operations $\cup$, $\setminus$, $\times$
Join $\bowtie$
Histograms
Equi-Width
Equi-Depth
Statistical Views

Database Profiles: IBM DB2

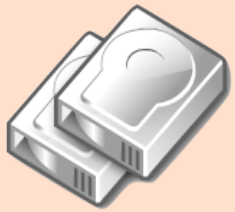
DB2 Excerpt of IBM DB2 catalog information for a TPC-H database

```
1  db2 => SELECT TABNAME, CARD, NPAGES
2  db2 (cont.) => FROM SYSCAT.TABLES
3  db2 (cont.) => WHERE TABSCHEMA = 'TPCH';

4  TABNAME         CARD                 NPAGES
5  --------------- -------------------- --------------------
6  ORDERS                       1500000                44331
7  CUSTOMER                      150000                 6747
8  NATION                            25                    2
9  REGION                             5                    1
10 PART                          200000                 7578
11 SUPPLIER                       10000                  406
12 PARTSUPP                      800000                31679
13 LINEITEM                     6001215               207888

14   8 record(s) selected.
```

- **Note**: Column CARD $\equiv |R|$, column NPAGES $\equiv N_R$.



Cardinality Estimation

Database Profiles

Assumptions

Estimating Operator Cardinality

Selection σ

Projection π

Set Operations $\cup$, $\setminus$, $\times$

Join $\bowtie$

Histograms

Equi-Width

Equi-Depth

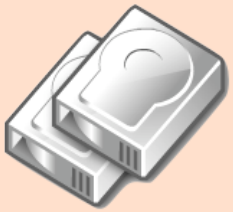
Statistical Views

Database Profile: Assumptions

In order to obtain tractable cardinality estimation formulae, assume *one of the following*:

Uniformity & independence (simple, yet rarely realistic)

All values of an attribute uniformly appear with the same probability. Values of different attributes are independent of each other.



Cardinality Estimation
Database Profiles
Assumptions
Estimating Operator Cardinality
Selection σ
Projection π
Set Operations $\cup$, $\setminus$, $\times$
Join $\bowtie$
Histograms
Equi-Width
Equi-Depth
Statistical Views

Database Profile: Assumptions

In order to obtain tractable cardinality estimation formulae, assume *one of the following*:

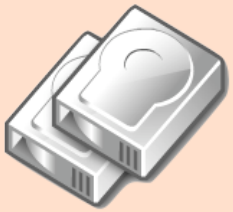
Uniformity & independence (simple, yet rarely realistic)

All values of an attribute uniformly appear with the same probability. Values of different attributes are independent of each other.

Worst case (unrealistic)

No knowledge about relation contents at all. In case of a selection σ_p, assume all records will satisfy predicate p.

(May only be used to compute upper bounds of expected cardinality.)



Cardinality Estimation
Database Profiles
Assumptions
Estimating Operator Cardinality
Selection σ
Projection π
Set Operations $\cup$, $\setminus$, $\times$
Join $\bowtie$
Histograms
Equi-Width
Equi-Depth
Statistical Views

Database Profile: Assumptions

In order to obtain tractable cardinality estimation formulae, assume *one of the following*:

Uniformity & independence (simple, yet rarely realistic)

All values of an attribute uniformly appear with the same probability. Values of different attributes are independent of each other.

Worst case (unrealistic)

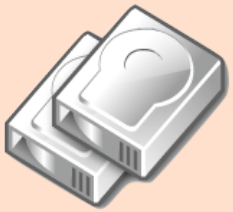
No knowledge about relation contents at all. In case of a selection σ_p, assume all records will satisfy predicate p.

(May only be used to compute upper bounds of expected cardinality.)

Perfect knowledge (unrealistic)

Details about the exact distribution of values are known. Requires huge catalog or prior knowledge of incoming queries.

(May only be used to compute lower bounds of expected cardinality.)



Cardinality Estimation
Database Profiles
Assumptions
Estimating Operator Cardinality
Selection σ
Projection π
Set Operations $\cup$, $\setminus$, $\times$
Join $\bowtie$
Histograms
Equi-Width
Equi-Depth
Statistical Views

Cardinality Estimation for σ (Equality Predicate)

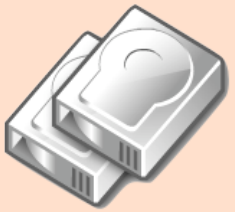
Query: $Q \equiv \sigma_{\mathtt{A}=c}(R)$

Selectivity $sel(\mathtt{A} = c)$	$1/V(\mathtt{A}, R)$	**Uniformity** 
Cardinality $\lvert Q \rvert$	$sel(\mathtt{A} = c) \cdot \lvert R \rvert$	
Record size $s(Q)$	$s(R)$	
Value Distribution $V(\mathtt{A}', Q)$	$\begin{cases} 1, & \text{for } \mathtt{A}' = \mathtt{A}, \\ c(V(\mathtt{A}', R), \lvert Q \rvert), & \text{otherwise.} \end{cases}$	

with (# of distinct colors obtained by drawing r balls from a bag of balls of m colors)[1]:

$$c(m, r) = \begin{cases} r, & \text{for } r < m/2, \\ (r + m)/3, & \text{for } m/2 \leqslant r < 2m, \\ m, & \text{for } r \geqslant 2m \end{cases}$$

[1] "Selection without replacement": $c(m, r) = m \cdot (1 - (1 - 1/m)^r)$.



Cardinality Estimation

Database Profiles

Assumptions

Estimating Operator Cardinality

Selection σ

Projection π

Set Operations $\cup$, $\setminus$, $\times$

Join $\bowtie$

Histograms

Equi-Width

Equi-Depth

Statistical Views

Selectivity Estimation for σ (Other Predicates)

- **Equality between attributes** ($Q \equiv \sigma_{\mathtt{A}=\mathtt{B}}(R)$):
  Approximate selectivity by

$$sel(\mathtt{A} = \mathtt{B}) \quad = \quad 1/\max(V(\mathtt{A}, R), V(\mathtt{B}, R)) \ .$$

  (Assumes that each value of the attribute with fewer distinct values has a corresponding match in the other attribute.)



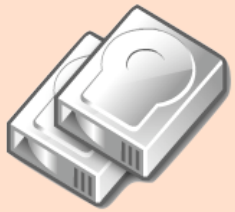
- **Range selections** ($Q = \sigma_{\mathtt{A}>c}(R)$):
  In the database profile, **maintain the minimum and maximum value** of attribute $\mathtt{A}$ in relation R, $Low(\mathtt{A}, R)$ and $High(\mathtt{A}, R)$.

  Approximate selectivity by



$$sel(\mathtt{A} > c) =$$

$$\begin{cases} \dfrac{High(\mathtt{A}, R) - c}{High(\mathtt{A}, R) - Low(\mathtt{A}, R)}, & Low(\mathtt{A}, R) \leqslant c \leqslant High(\mathtt{A}, R) \\ 0, & \text{otherwise} \end{cases}$$



Cardinality Estimation

Database Profiles

Assumptions

Estimating Operator Cardinality

Selection σ

Projection π

Set Operations $\cup$, $\setminus$, $\times$

Join $\bowtie$

Histograms

Equi-Width

Equi-Depth

Statistical Views

Cardinality Estimation for π

- For $Q \equiv \pi_L(R)$, estimating the number of result rows is difficult ($L = \langle \mathtt{A}_1, \mathtt{A}_2, \ldots, \mathtt{A}_n \rangle$: list of projection attributes):

$Q \equiv \pi_L(R)$

Cardinality $|Q|$

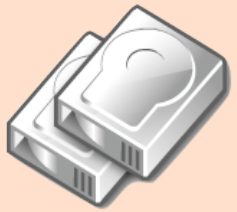
$$\begin{cases} V(\mathtt{A}, R), & \text{if } L = \langle \mathtt{A} \rangle \\ |R|, & \text{if keys of } R \in L \\ |R|, & \text{no dup. elim.} \\ \min\left(|R|, \prod_{\mathtt{A}_i \in L} V(\mathtt{A}_i, R)\right), & \text{otherwise} \end{cases}$$

Independence



Record size $s(Q)$ $\quad \sum_{\mathtt{A}_i \in L} s(\mathtt{A}_i)$

Val. Dist. $V(\mathtt{A}_i, Q)$ $\quad V(\mathtt{A}_i, R)$ for $\mathtt{A}_i \in L$



Cardinality Estimation

Database Profiles

Assumptions

Estimating Operator Cardinality

Selection σ

Projection π

Set Operations $\cup$, $\setminus$, $\times$

Join $\bowtie$

Histograms

Equi-Width

Equi-Depth

Statistical Views

Cardinality Estimation for $\cup, \setminus, \times$

$Q \equiv R \cup S$

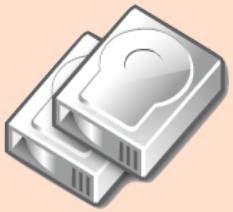
$$\begin{aligned} |Q| &\leqslant |R| + |S| \\ s(Q) &= s(R) = s(S) \qquad \text{schemas of } R,S \text{ identical} \\ V(\mathtt{A}, Q) &\leqslant V(\mathtt{A}, R) + V(\mathtt{A}, S) \end{aligned}$$

$Q \equiv R \setminus S$

$$\begin{aligned} \max(0, |R| - |S|) \leqslant |Q| &\leqslant |R| \\ s(Q) &= s(R) = s(S) \\ V(\mathtt{A}, Q) &\leqslant V(\mathtt{A}, R) \end{aligned}$$

$Q \equiv R \times S$

$$\begin{aligned} |Q| &= |R| \cdot |S| \\ s(Q) &= s(R) + s(S) \\ V(\mathtt{A}, Q) &= \begin{cases} V(\mathtt{A}, R), & \text{for } \mathtt{A} \in R \\ V(\mathtt{A}, S), & \text{for } \mathtt{A} \in S \end{cases} \end{aligned}$$



Cardinality Estimation
Database Profiles
Assumptions
Estimating Operator Cardinality
Selection σ
Projection π
Set Operations $\cup, \setminus, \times$
Join $\bowtie$
Histograms
Equi-Width
Equi-Depth
Statistical Views

Cardinality Estimation for $\bowtie$

- Cardinality estimation for the general join case is challenging.
- A special, yet very common case: **foreign-key relationship** between input relations R and S:

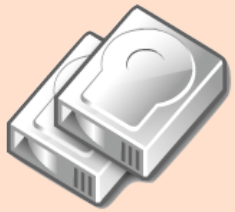
Establish a foreign key relationship (SQL)

```
1 CREATE TABLE R (A INTEGER NOT NULL,
2                 ...
3                 PRIMARY KEY (A));
4 CREATE TABLE S (...,
5                 A INTEGER NOT NULL,
6                 ...
7                 FOREIGN KEY (A) REFERENCES R);
```

$Q \equiv R \bowtie_{R.A=S.A} S$

The foreign key constraint guarantees $\pi_A(S) \subseteq \pi_A(R)$. Thus:

$$|Q| \quad = \quad |S| \ .$$



Cardinality Estimation
Database Profiles
Assumptions
Estimating Operator Cardinality
Selection σ
Projection π
Set Operations $\cup$, $\setminus$, $\times$
Join $\bowtie$
Histograms
Equi-Width
Equi-Depth
Statistical Views

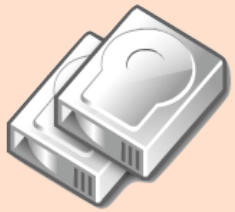
Cardinality Estimation for $\bowtie$

$Q \equiv R \bowtie_{R.\mathtt{A}=S.\mathtt{B}} S$

$$|Q| = \begin{cases} \dfrac{|R| \cdot |S|}{V(\mathtt{A}, R)}, & \pi_{\mathtt{B}}(S) \subseteq \pi_{\mathtt{A}}(R) \\ \dfrac{|R| \cdot |S|}{V(\mathtt{B}, S)}, & \pi_{\mathtt{A}}(R) \subseteq \pi_{\mathtt{B}}(S) \end{cases}$$

$$s(Q) = s(R) + s(S)$$

$$V(\mathtt{A}', Q) \leqslant \begin{cases} V(\mathtt{A}', R), & \text{if } \mathtt{A}' \text{ attribute in } R \\ V(\mathtt{A}', S), & \text{otherwise} \end{cases}$$

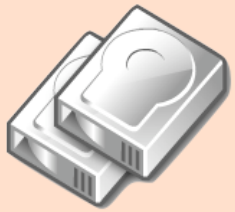


Cardinality Estimation
Database Profiles
Assumptions
Estimating Operator Cardinality
Selection σ
Projection π
Set Operations $\cup$, $\setminus$, $\times$
Join $\bowtie$
Histograms
Equi-Width
Equi-Depth
Statistical Views

Histograms

- In realistic database instances, values are **not uniformly distributed** in an attribute's **active domain** (actual values found in a column).
- To keep track of this non-uniformity for an attribute A, maintain a **histogram** to **approximate the actual distribution**:
  1. Divide the active domain of A into adjacent intervals by selecting **boundary values** b_i.
  2. Collect statistical parameters for each interval between boundaries, *e.g.*,
     - # of rows r with $b_{i-1} < r.\mathtt{A} \leqslant b_i$, or
     - # of distinct A values in interval $(b_{i-1}, b_i]$.
- The histogram intervals are also referred to as **buckets.**

(↗ Y. Ioannidis: The History of Histograms (Abridged), *Proc. VLDB 2003*)



Cardinality Estimation
Database Profiles
Assumptions
Estimating Operator Cardinality
Selection σ
Projection π
Set Operations $\cup$, $\setminus$, $\times$
Join $\bowtie$
Histograms
Equi-Width
Equi-Depth
Statistical Views

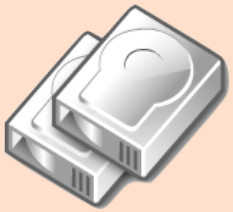
Histograms in IBM DB2

DB2 Histogram maintained for a column in a TPC-H database

```
1  SELECT SEQNO, COLVALUE, VALCOUNT
2    FROM SYSCAT.COLDIST
3   WHERE TABNAME = 'LINEITEM'
4     AND COLNAME = 'L_EXTENDEDPRICE'
5     AND TYPE = 'Q';

6  SEQNO COLVALUE                VALCOUNT
7  ----- ---------------------- --------
8      1 +0000000000996.01          3001
9      2 +0000000004513.26        315064
10     3 +0000000007367.60        633128
11     4 +0000000011861.82        948192
12     5 +0000000015921.28       1263256
13     6 +0000000019922.76       1578320
14     7 +0000000024103.20       1896384
15     8 +0000000027733.58       2211448
16     9 +0000000031961.80       2526512
17    10 +0000000035584.72       2841576
18    11 +0000000039772.92       3159640
19    12 +0000000043395.75       3474704
20    13 +0000000047013.98       3789768
21                   ⋮
```

- Catalog table `SYSCAT.COLDIST` also contains information like
  - the n most frequent values (and their frequency),
  - the number of distinct values in each bucket.
- Histograms may even be manipulated **manually** to tweak optimizer decisions.



Cardinality Estimation

Database Profiles

Assumptions

Estimating Operator Cardinality

Selection σ

Projection π

Set Operations $\cup$, $\setminus$, $\times$

Join $\bowtie$

Histograms

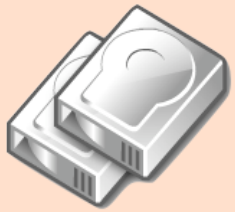
Equi-Width

Equi-Depth

Statistical Views

Histograms

- Two types of histograms are widely used:
  1. **Equi-Width Histograms.**
     All buckets have the **same width**, *i.e.*, boundary $b_i = b_{i-1} + w$, for some fixed w.
  2. **Equi-Depth Histograms.**
     All buckets contain the **same number of rows** (*i.e.*, their width is varying).
- Equi-depth histograms (2) are able to adapt to data skew (high uniformity).
- The number of buckets is the *tuning knob* that defines the **tradeoff** between estimation quality (**histogram resolution**) and **histogram size**: catalog space is limited.

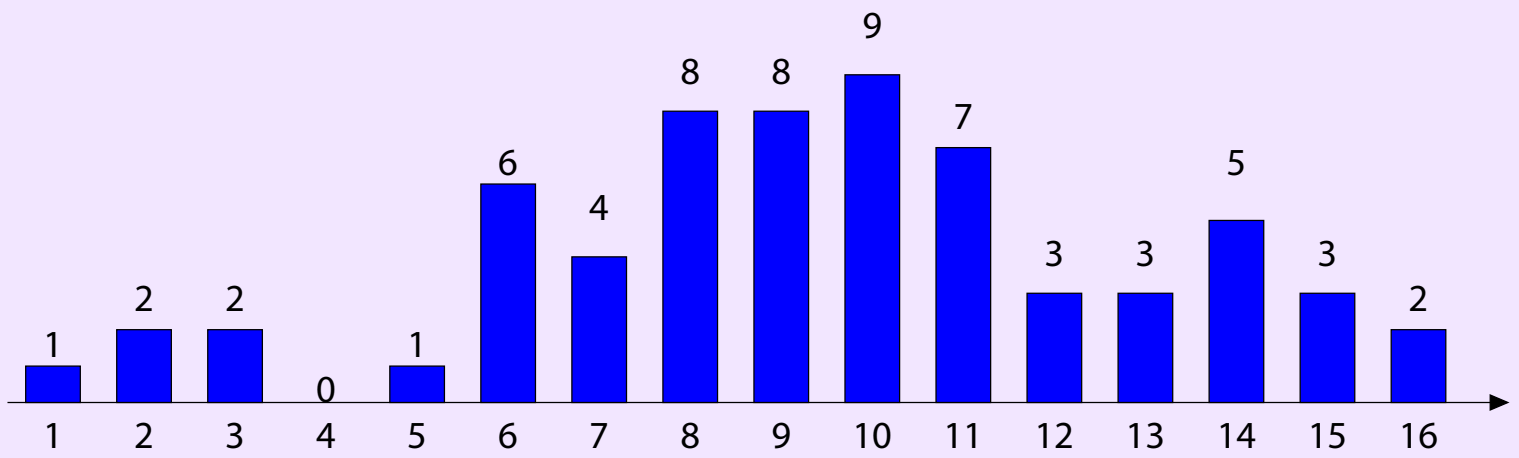


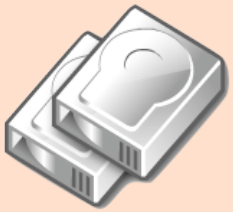
Cardinality Estimation
Database Profiles
Assumptions
Estimating Operator Cardinality
Selection σ
Projection π
Set Operations $\cup$, $\setminus$, $\times$
Join $\bowtie$
Histograms
Equi-Width
Equi-Depth
Statistical Views

Equi-Width Histograms

Example (Actual value distribution)

Column `A` of SQL type `INTEGER` (domain $\{\ldots, -2, -1, 0, 1, 2, \ldots\}$).
Actual non-uniform distribution in relation R:





Cardinality Estimation

Database Profiles

Assumptions

Estimating Operator Cardinality

Selection σ

Projection π

Set Operations $\cup$, $\setminus$, $\times$

Join $\bowtie$

Histograms

Equi-Width

Equi-Depth

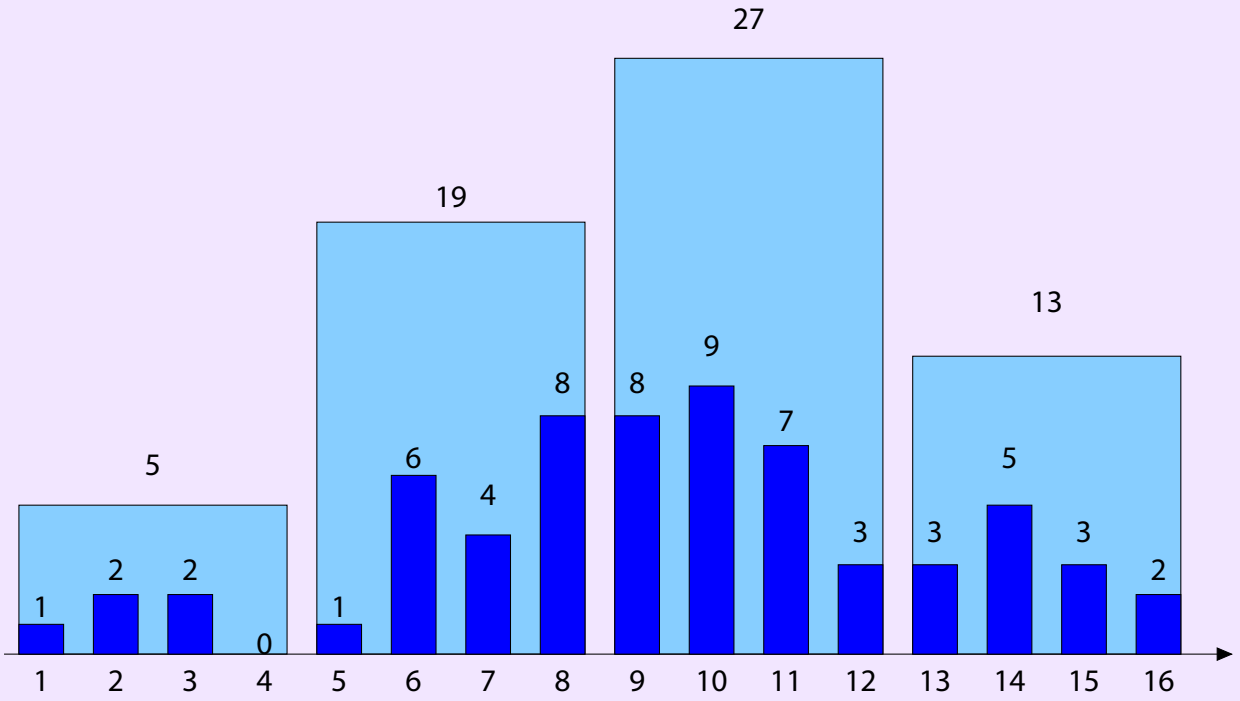
Statistical Views

Equi-Width Histograms

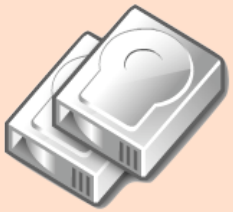
- Divide **active domain** of attribute A into B buckets of equal width. The **bucket width** w will be

$$w = \frac{High(\mathtt{A}, R) - Low(\mathtt{A}, R) + 1}{B}$$

Example (Equi-width histogram ($B = 4$))



- Maintain **sum of value frequencies** in each bucket (in addition to bucket boundaries b_i).



Cardinality Estimation

Database Profiles

Assumptions

Estimating Operator Cardinality

Selection σ

Projection π

Set Operations $\cup$, $\setminus$, $\times$

Join $\bowtie$

Histograms

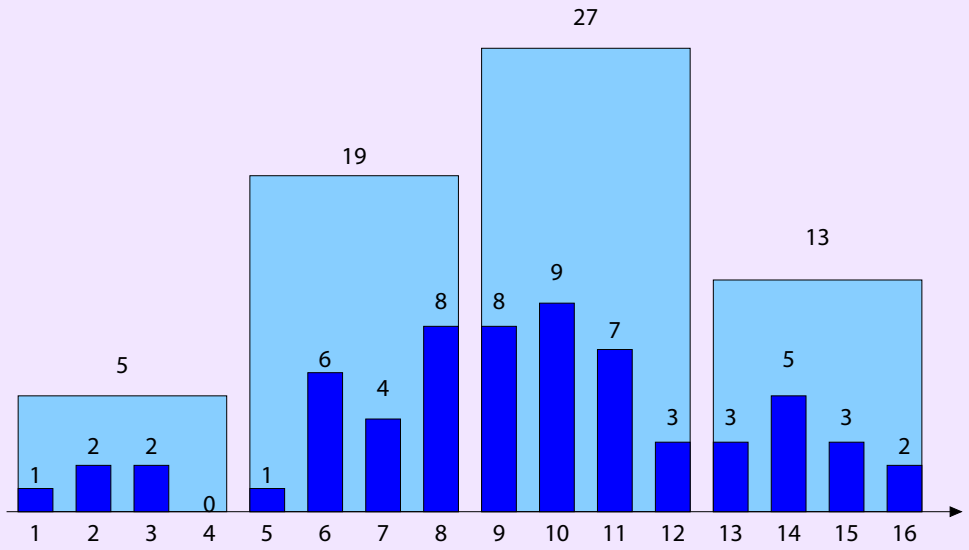
Equi-Width

Equi-Depth

Statistical Views

Equi-Width Histograms: Equality Selections

Example ($Q \equiv \sigma_{A=5}(R)$)

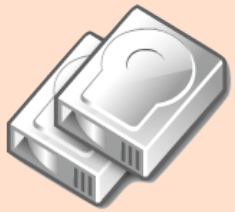


- Value 5 is in bucket $[5, 8]$ (with 19 tuples)
- Assume **uniform distribution within the bucket**:

$$|Q| = 19/w = 19/4 \approx 5 \ .$$

Actual: $|Q| = 1$

What would be the cardinality under the uniformity assumption (no histogram)?



Cardinality Estimation

Database Profiles

Assumptions

Estimating Operator Cardinality

Selection σ

Projection π

Set Operations $\cup$, $\setminus$, $\times$

Join $\bowtie$

Histograms

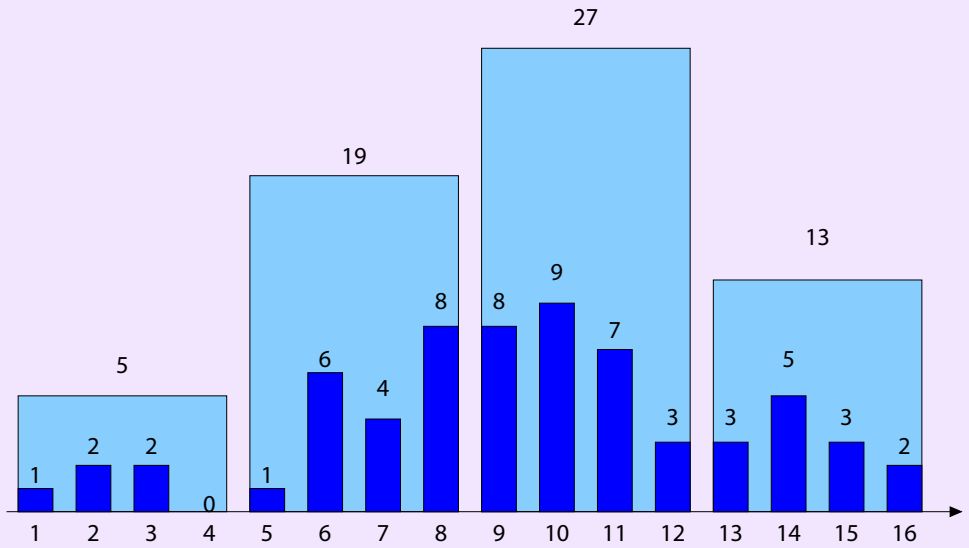
Equi-Width

Equi-Depth

Statistical Views

Equi-Width Histograms: Range Selections

Example ($Q \equiv \sigma_{\texttt{A}>7 \text{ AND } \texttt{A}\leqslant 16}(R)$)

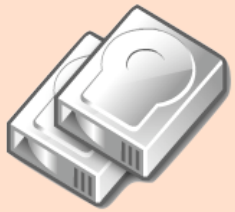


- Query interval $(7, 16]$ covers buckets $[9, 12]$ and $[13, 16]$. Query interval touches $[5, 8]$.

$$|Q| = 27 + 13 + {}^{19}/_{4} \approx 45 \quad .$$

Actual: $|Q| = 48$

What would be the cardinality under the uniformity assumption (no histogram)?



Cardinality Estimation

Database Profiles

Assumptions

Estimating Operator Cardinality

Selection σ

Projection π

Set Operations $\cup$, $\setminus$, $\times$

Join $\bowtie$

Histograms

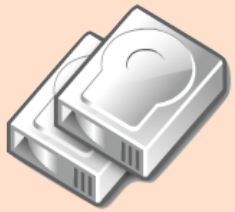
Equi-Width

Equi-Depth

Statistical Views

Equi-Width Histogram: Construction

- To **construct** an equi-width histogram for relation R, attribute $\mathtt{A}$:
  1. Compute boundaries b_i from $High(\mathtt{A}, R)$ and $Low(\mathtt{A}, R)$.
  2. Scan R once sequentially.
  3. While scanning, maintain B running tuple frequency counters, one for each bucket.
- If scanning R in step 2 is prohibitive, scan small sample $R_{sample} \subset R$, then scale frequency counters by $|R|/|R_{sample}|$.
- To **maintain** the histogram under insertions (deletions):
  1. Simply increment (decremeent) frequency counter in affected bucket.



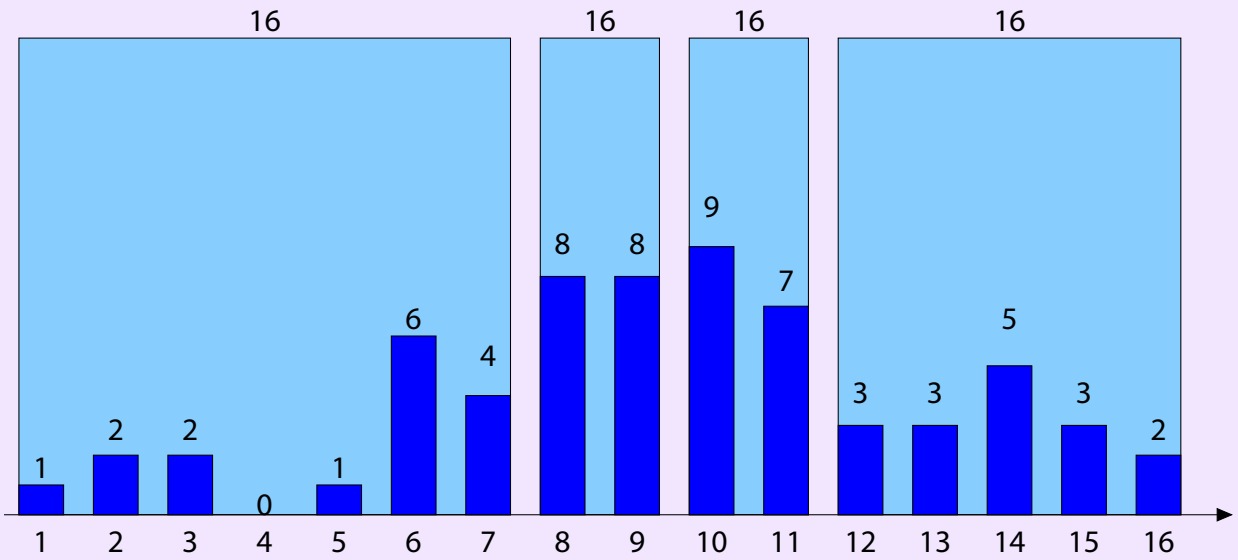
Cardinality Estimation
Database Profiles
Assumptions
Estimating Operator Cardinality
Selection σ
Projection π
Set Operations $\cup$, $\setminus$, $\times$
Join $\bowtie$
Histograms
Equi-Width
Equi-Depth
Statistical Views

Equi-Depth Histograms

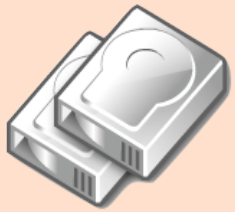
- Divide **active domain** of attribute A into B buckets of roughly the same number of tuples in each bucket, **depth** d of each bucket will be

$$d = \frac{|R|}{B} \quad .$$

Example (Equi-depth histogram ($B = 4, d = 16$))



- Maintain **depth** (and bucket boundaries b_i).



Cardinality Estimation

Database Profiles

Assumptions

Estimating Operator Cardinality

Selection σ

Projection π

Set Operations $\cup$, $\setminus$, $\times$

Join $\bowtie$

Histograms

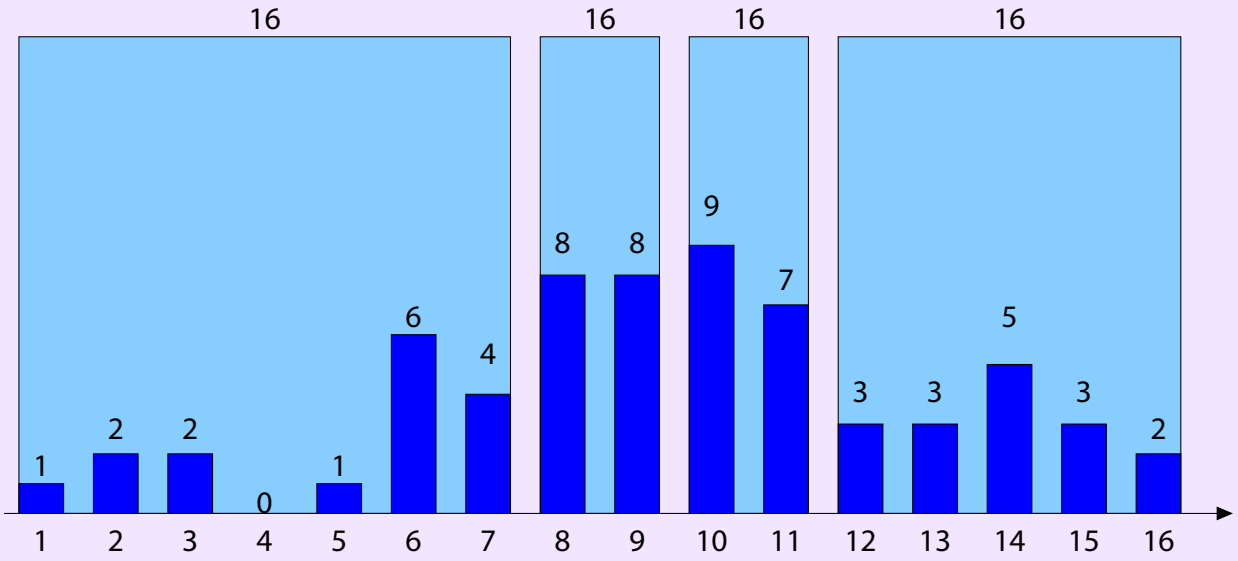
Equi-Width

Equi-Depth

Statistical Views

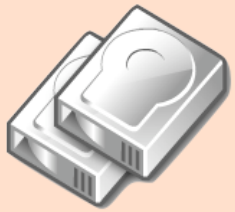
Equi-Depth Histograms

Example (Equi-depth histogram ($B = 4, d = 16$))



Intuition:

- High value frequencies are more important than low value frequencies.
- Resolution of histogram adapts to skewed value distributions.



Cardinality Estimation

Database Profiles

Assumptions

Estimating Operator Cardinality

Selection σ

Projection π

Set Operations $\cup$, $\setminus$, $\times$

Join $\bowtie$

Histograms

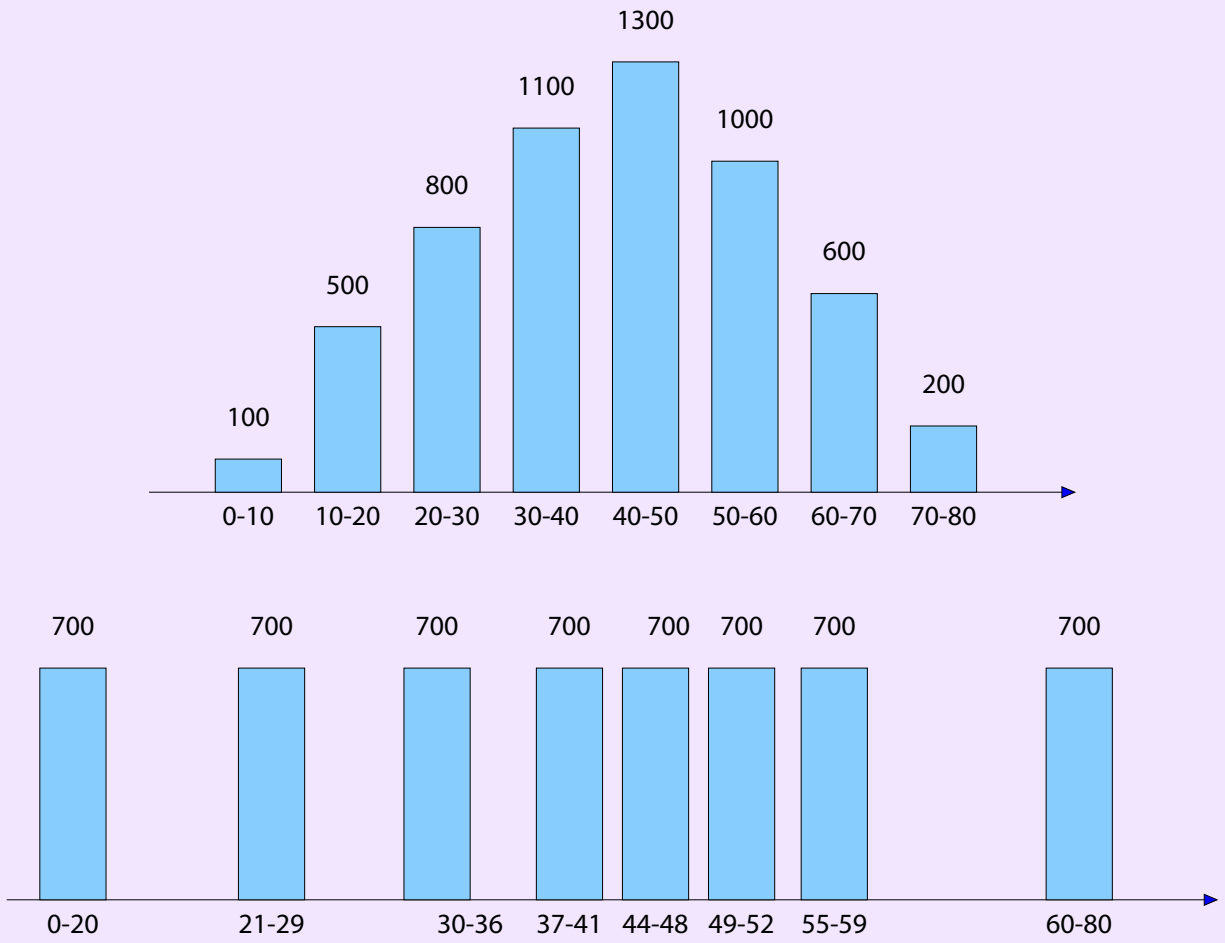
Equi-Width

Equi-Depth

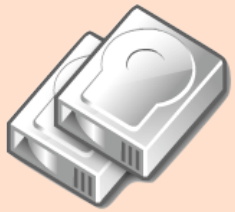
Statistical Views

Equi-Width vs. Equi-Depth Histograms

Example (Histogram on *customer age* attribute $(B = 8, |R| = 5{,}600)$)



- Equi-depth histogram "invests" bytes in the densely populated customer age region between 30 and 59.



Cardinality Estimation

Database Profiles

Assumptions

Estimating Operator Cardinality

Selection σ

Projection π

Set Operations $\cup$, $\setminus$, $\times$

Join $\bowtie$

Histograms

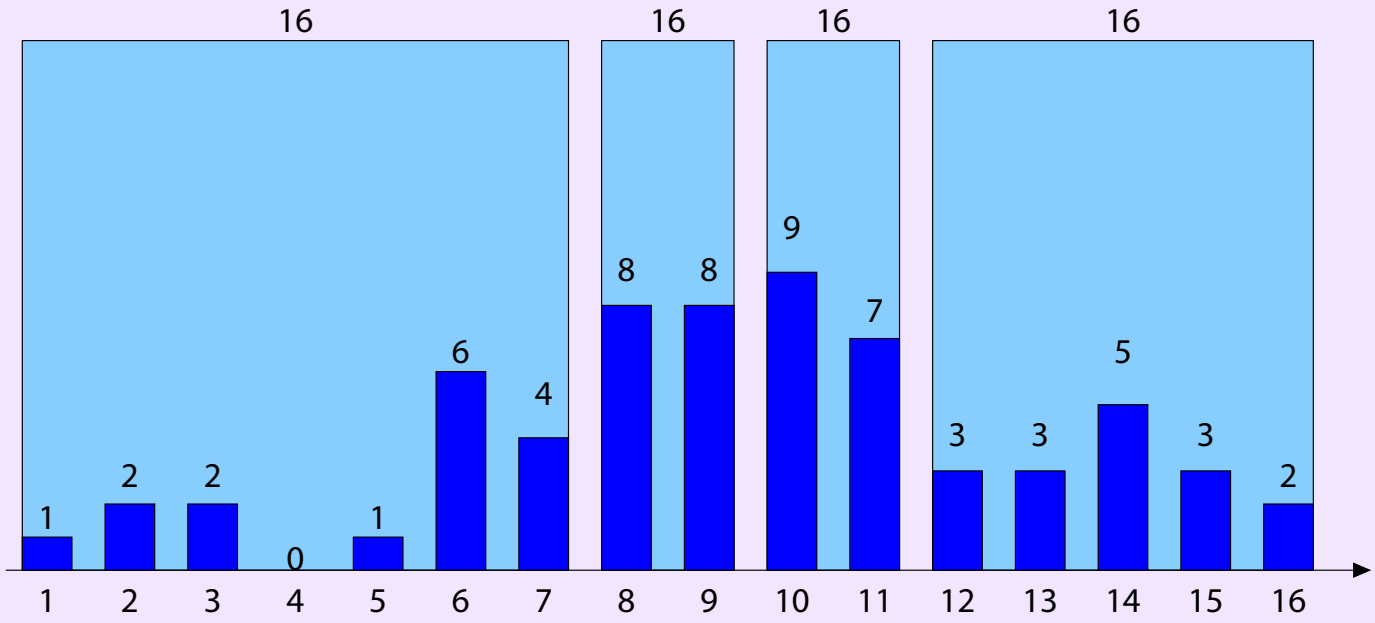
Equi-Width

Equi-Depth

Statistical Views

Equi-Depth Histograms: Equality Selections

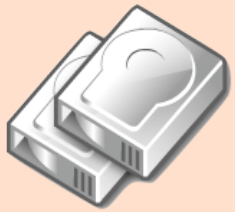
Example ($Q \equiv \sigma_{A=5}(R)$)



- Value 5 is in first bucket $[1, 7]$ (with $d = 16$ tuples)
- Assume **uniform distribution within the bucket**:

$$|Q| = d/7 = 16/7 \approx 2 \ .$$

(Actual: $|Q| = 1$)



Cardinality Estimation

Database Profiles

Assumptions

Estimating Operator Cardinality

Selection σ

Projection π

Set Operations $\cup$, $\setminus$, $\times$

Join $\bowtie$

Histograms

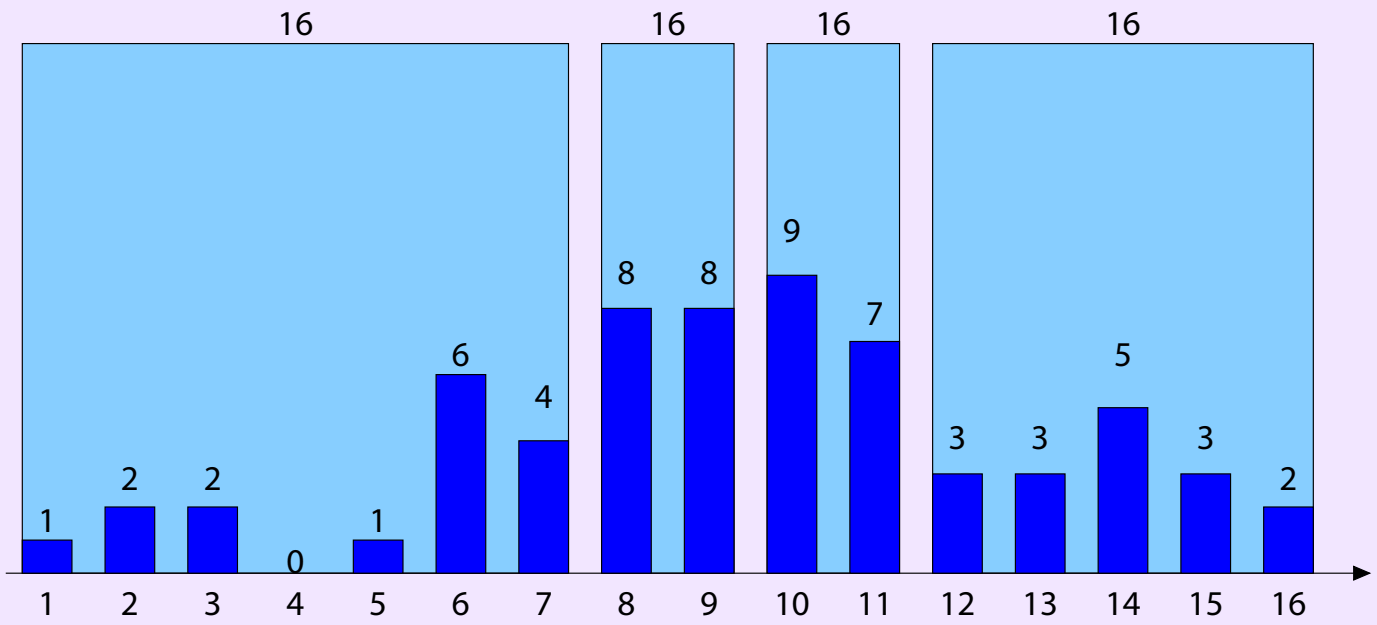
Equi-Width

Equi-Depth

Statistical Views

Equi-Depth Histograms: Range Selections

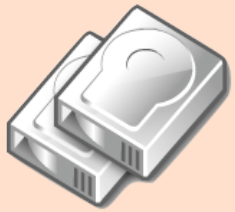
Example ($Q \equiv \sigma_{\texttt{A}>5 \texttt{ AND } \texttt{A}\leqslant 16}(R)$)



- Query interval $(5, 16]$ covers buckets $[8, 9]$, $[10, 11]$ and $[12, 16]$ (all with $d = 16$ tuples). Query interval touches $[1, 7]$.

$$|Q| = 16 + 16 + 16 + {}^{2}/_{7} \cdot 16 \approx 53 \quad .$$

(Actual: $|Q| = 59$)



Cardinality Estimation

Database Profiles

Assumptions

Estimating Operator Cardinality

Selection σ

Projection π

Set Operations $\cup$, $\setminus$, $\times$

Join $\bowtie$

Histograms

Equi-Width

Equi-Depth

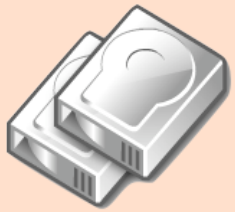
Statistical Views

Equi-Depth Histograms: Construction

- To **construct** an equi-depth histogram for relation R, attribute $\mathtt{A}$:
  1. Compute depth $d = |R|/B$.
  2. Sort R by sort criterion $\mathtt{A}$.
  3. $b_0 = Low(\mathtt{A}, R)$, then determine the b_i by dividing the sorted R into chunks of size d.

Example ($B = 4, |R| = 64$)

1. $d = {}^{64}/_{4} = 16$.
2. Sorted $R.\mathtt{A}$:
   $\langle 1,2,2,3,3,5,6,6,6,6,6,6,7,7,7,7,8,8,8,8,8,8,8,8,9,9,9,9,9,9,9,9,10,10,\ldots\rangle$
3. Boundaries of d-sized chunks in sorted R:
   $\langle \underbrace{1,2,2,3,3,5,6,6,6,6,6,6,7,7,7,7}_{b_1=7}, \underbrace{8,8,8,8,8,8,8,8,9,9,9,9,9,9,9,9}_{b_2=9}, 10,10,\ldots\rangle$



Cardinality Estimation
Database Profiles
Assumptions
Estimating Operator Cardinality
Selection σ
Projection π
Set Operations $\cup$, $\setminus$, $\times$
Join $\bowtie$
Histograms
Equi-Width
Equi-Depth
Statistical Views

A Cardinality (Mis-)Estimation Scenario

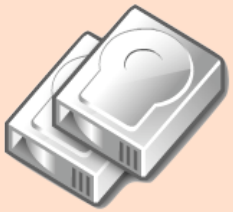
- Because exact cardinalities and estimated selectivity information is provided for base tables only, the DBMS relies on **projected cardinalities** for derived tables.
- In the case of foreign key joins, IBM DB2 promotes selectivity factors for one join input to the join result, for example.

Example (Selectivity promotion; $\mathtt{K}$ is key of S, $\pi_{\mathtt{A}}(R) \subseteq \pi_{\mathtt{K}}(S)$)

$$R \bowtie_{R.\mathtt{A}=S.\mathtt{K}} (\sigma_{\mathtt{B}=10}(S))$$

If $sel(\mathtt{B} = 10) = x$, then assume that the join will yield $x \cdot |R|$ rows.

- Whenever the value distribution of $\mathtt{A}$ in R does not match the distribution of $\mathtt{B}$ in S, the cardinality estimate may be severly off.



Cardinality Estimation
Database Profiles
Assumptions
Estimating Operator Cardinality
Selection σ
Projection π
Set Operations $\cup$, $\setminus$, $\times$
Join $\bowtie$
Histograms
Equi-Width
Equi-Depth
Statistical Views

A Cardinality (Mis-)Estimation Scenario

Example (Excerpt of a data warehouse)

Dimension table STORE:

STOREKEY	STORE_NUMBER	CITY	STATE	DISTRICT
...	...	...	...	...

} 63 rows

Dimension table PROMOTION:

PROMOKEY	PROMOTYPE	PROMODESC	PROMOVALUE
...	...	...	...

} 35 rows

Fact table DAILY_SALES:

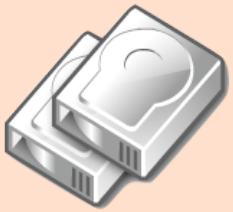
STOREKEY	CUSTKEY	PROMOKEY	SALES_PRICE
...	...	...	...

} 754 069 426 rows

Let the tables be arranged in a **star schema**:

- The **fact table** references the **dimension tables**,
- the dimension tables are small/stable, the fact table is large/continously update on each sale.

$\Rightarrow$ Histograms are maintained for the dimension tables.



Cardinality Estimation
Database Profiles
Assumptions
Estimating Operator Cardinality
Selection σ
Projection π
Set Operations $\cup$, $\setminus$, $\times$
Join $\bowtie$
Histograms
Equi-Width
Equi-Depth
Statistical Views

A Cardinality (Mis-)Estimation Scenario

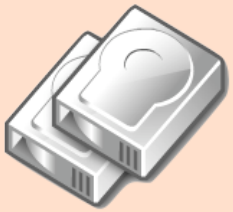
DB2 Query against the data warehouse

Find the number of those sales in store '01' (18 of the overall 63 locations) that were the result of the sales promotion of type 'XMAS' (**"star join"**):

```
1   SELECT COUNT(*)
2   FROM   STORE d1, PROMOTION d2, DAILY_SALES f
3   WHERE  d1.STOREKEY = f.STOREKEY
4   AND    d2.PROMOKEY = f.PROMOKEY
5   AND    d1.STORE_NUMBER = '01'
6   AND    d2.PROMOTYPE = 'XMAS'
```

The query yields 12,889,514 rows. The histograms lead to the following selectivity estimates:

$$sel(\texttt{STORE_NUMBER = '01'}) = {}^{18}/_{63} \quad (28.57\,\%)$$
$$sel(\texttt{PROMOTYPE = 'XMAS'}) = {}^{1}/_{35} \quad (2.86\,\%)$$



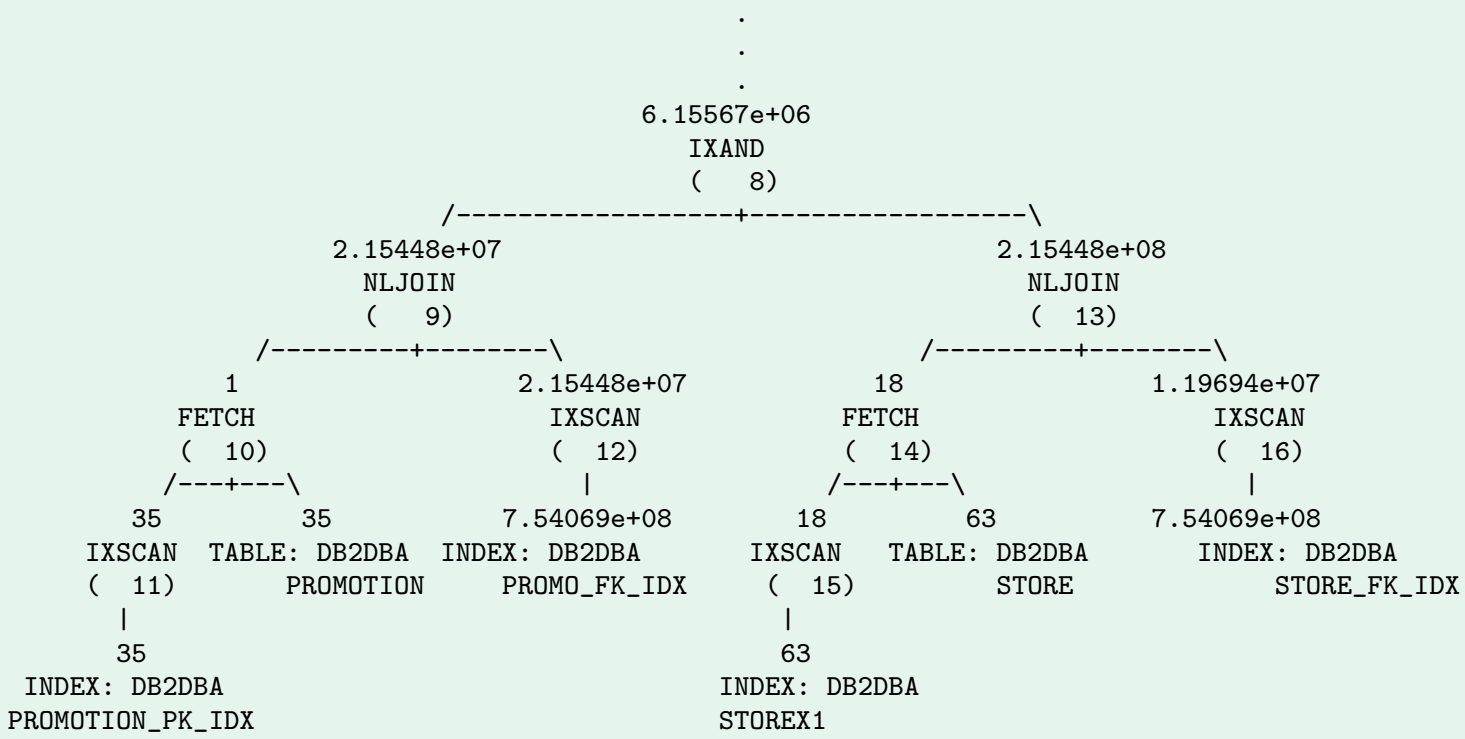
Cardinality Estimation
Database Profiles
Assumptions
Estimating Operator Cardinality
Selection σ
Projection π
Set Operations $\cup$, $\setminus$, $\times$
Join $\bowtie$
Histograms
Equi-Width
Equi-Depth
Statistical Views

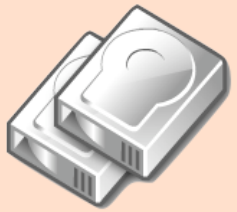
A Cardinality (Mis-)Estimation Scenario

DB2 Estimated cardinalities and selected plan

```
1   SELECT COUNT(*)
2   FROM   STORE d1, PROMOTION d2, DAILY_SALES f
3   WHERE  d1.STOREKEY = f.STOREKEY
4   AND    d2.PROMOKEY = f.PROMOKEY
5   AND    d1.STORE_NUMBER = '01'
6   AND    d2.PROMOTYPE = 'XMAS'
```

Plan fragment (top numbers indicates estimated cardinality):





Cardinality Estimation

Database Profiles

Assumptions

Estimating Operator Cardinality

Selection σ

Projection π

Set Operations $\cup$, $\setminus$, $\times$

Join $\bowtie$

Histograms

Equi-Width

Equi-Depth

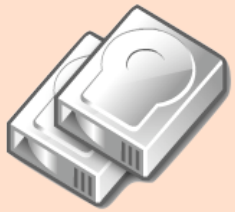
Statistical Views

IBM DB2: Statistical Views

- To provide database profile information (estimate cardinalities, value distributions, . . .) for **derived tables**:
  1. Define a **view** that precomputes the derived table (or possibly a small sample of it, IBM DB2: 10 %),
  2. use the view result to gather and keep **statistics**, then delete the result.

DB2 Statistical views

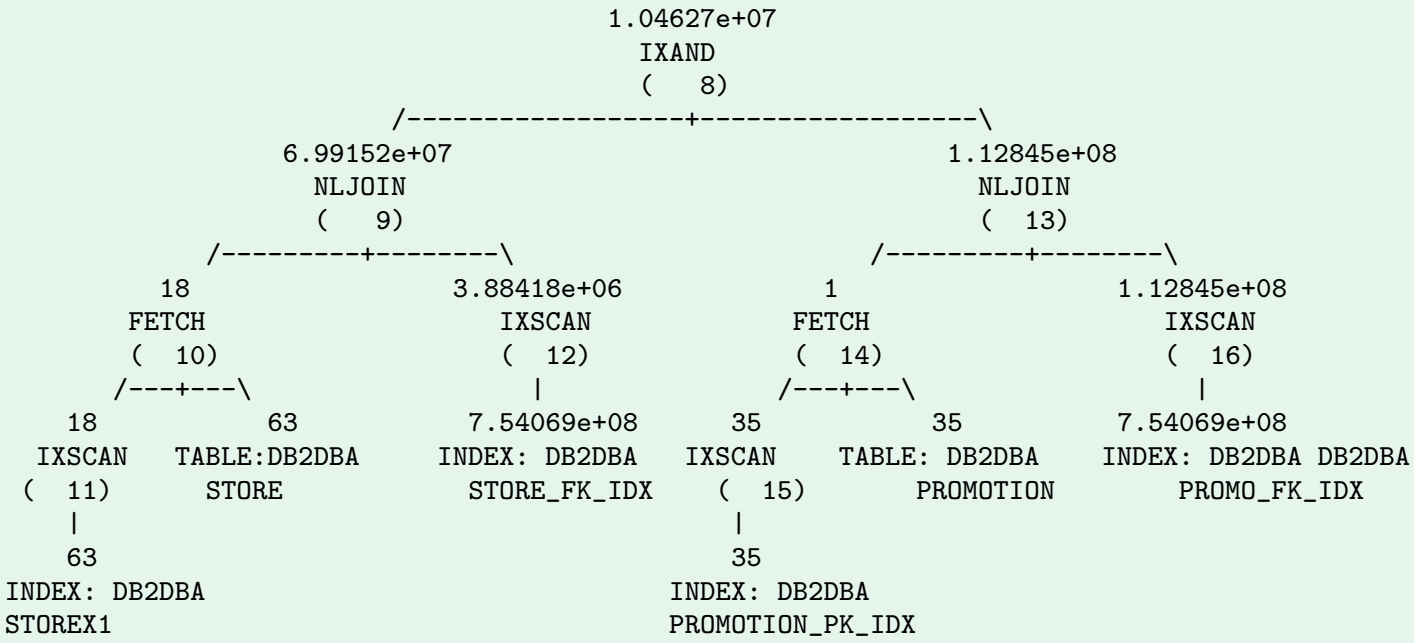
```
1  CREATE VIEW sv_store_dailysales AS
2     (SELECT s.*
3      FROM   STORE s, DAILY_SALES ds
4      WHERE  s.STOREKEY = ds.STOREKEY);
5
6  CREATE VIEW sv_promotion_dailysales AS
7     (SELECT p.*
8      FROM   PROMOTION p, DAILY_SALES ds
9      WHERE  p.PROMOKEY = ds.PROMOKEY);
10
11 ALTER VIEW sv_store_dailysales ENABLE QUERY OPTIMIZATION;
12 ALTER VIEW sv_promotion_dailysales ENABLE QUERY OPTIMIZATION;
13
14 RUNSTATS ON TABLE sv_store_dailysales WITH DISTRIBUTION;
15 RUNSTATS ON TABLE sv_promotion_dailysales WITH DISTRIBUTION;
```

Cardinality Estimation
Database Profiles
Assumptions
Estimating Operator Cardinality
Selection σ
Projection π
Set Operations $\cup$, $\setminus$, $\times$
Join $\bowtie$
Histograms
Equi-Width
Equi-Depth
Statistical Views

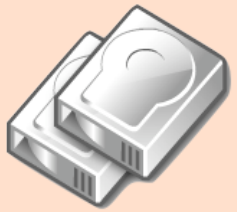
Cardinality Estimation with Statistical Views

DB2 Estimated cardinalities and selected plan after reoptimization



Note new estimated selectivities *after* join:

- Selectivity of `PROMOTYPE = 'XMAS'` now only 14.96 % (was: 2.86 %)
- Selectivity of `STORE_NUMBER = '01'` now 9.27 % (was: 28.57 %)



Cardinality Estimation

Database Profiles

Assumptions

Estimating Operator Cardinality

Selection σ

Projection π

Set Operations $\cup$, $\setminus$, $\times$

Join $\bowtie$

Histograms

Equi-Width

Equi-Depth

Statistical Views