

## Approximation Techniques for Data Management Systems

"We are drowning in data but starved for knowledge"

John Naibitt



## Traditional Query Processing



- Exact answers **NOT** always required
  - DSS applications usually *exploratory*: early feedback to help identify "interesting" regions
  - Aggregate queries*: precision to "last decimal" not needed
    - e.g., "What percentage of the US sales are in NJ?"

2

## Fast Approximate Answers


- Primarily for Aggregate queries
- Goal is to quickly report the leading digits of answers
  - In **seconds** instead of minutes or hours
  - Most useful if can provide **error guarantees**

E.g., Average salary  
 \$59,000 +/- \$500 (with 95% confidence) in 10 seconds  
 vs. \$59,152.25 in 10 minutes

- Achieved by answering the query based on *compact synopses* of the data
- Speed-up obtained because synopses are **orders of magnitude smaller** than the original data

3

## Approximate Query Processing



- How do you **build effective data synopses**???

4

## Sampling: Basics

- Idea: A small random sample  $S$  of the data often well-represents all the data
  - For a fast approx answer, apply the query to  $S$  & "scale" the result
  - E.g.,  $R.A$  is  $\{0,1\}$ ,  $S$  is a 20% sample
 

1 1 0 1  
 1 1 1 0 0 0  
 0 1 1 1 1 0 1  
 1 1 0 1 0 1 1  
 0 1 1 0

**R.A**  
 Red =  
 in  $S$

select count(\*) from  $R$  where  $R.A = 0$   
 ↓  
 select 5 \* count(\*) from  $S$  where  $S.A = 0$   
 Est. count =  $5 * 2 = 10$ , Exact count = 10
- Unbiased**: For expressions involving count, sum, avg: the estimator is unbiased, i.e., the expected value of the answer is the actual answer, even for (most) queries with predicates!
- Leverage extensive literature on **confidence intervals** for sampling
  - Actual answer is within the interval  $[a,b]$  with a given probability
  - E.g.,  $54,000 \pm 600$  with prob  $\geq 90\%$

5

## Sampling: Confidence Intervals

Method	90% Confidence Interval ( $\pm$ )	Guarantees?
Central Limit Theorem	$1.65 * \sigma(S) / \sqrt{( S )}$	as $ S  \rightarrow \infty$
Hoeffding	$1.22 * (MAX-MIN) / \sqrt{( S )}$	always
Chebyshev (known $\sigma(R)$ )	$3.16 * \sigma(R) / \sqrt{( S )}$	always
Chebyshev (est. $\sigma(R)$ )	$3.16 * \sigma(S) / \sqrt{( S )}$	as $\sigma(S) \rightarrow \sigma(R)$

**Confidence intervals for Average**: select avg( $R.A$ ) from  $R$   
 (Can replace  $R.A$  with any arithmetic expression on the attributes in  $R$ )  
 $\sigma(R)$  = standard deviation of the values of  $R.A$ ;  $\sigma(S)$  = s.d. for  $S.A$

- If predicates,  $S$  above is subset of sample that satisfies the predicate
- Quality of the estimate depends only on the **variance in  $R$  &  $|S|$  after the predicate**: So 10K sample may suffice for 10B row relation!
  - Advantage of larger samples: can handle more selective predicates

6

## Sampling from Databases

- Sampling disk-resident data is slow
  - Row-level sampling has **high I/O cost**:
    - must bring in entire disk block to get the row
  - Block-level sampling: rows may be **highly correlated**
  - Random access pattern**, possibly via an index
  - Need to account for the variable number of rows in a page, children in an index node, etc.
- Alternatives
  - Random physical clustering**: destroys "natural" clustering
  - Precomputed samples**: must incrementally maintain (at specified size)
    - Fast to use: packed in disk blocks, can sequentially scan, can store as relation and leverage full DBMS query support, can store in main memory

7

## One-Pass Uniform Sampling

- Best choice for incremental maintenance
  - Low overheads, no random data access
- Reservoir Sampling [Vit85]: **Maintains a sample  $S$  of a fixed-size  $M$** 
  - Add each new item to  $S$  with probability  $M/N$ , where  $N$  is the current number of data items
  - If add an item, evict a random item from  $S$
  - Instead of flipping a coin for each item, determine the number of items to skip before the next to be added to  $S$


8

## Histograms

- Partition attribute value(s) domain into a set of buckets
- Issues:
  - How to partition
  - What to store for each bucket
  - How to estimate an answer using the histogram
- Long history of use for selectivity estimation within a query optimizer
- Recently explored as a tool for fast approximate query processing

9

## 1-D Histograms



- Number of buckets  $B \ll$  domain size
- Each bucket just stores a **total count**
  - Distributed uniformly across values in the bucket
- Partition criteria**
  - Equi-width**: equal number of domain values per bucket (bad!!)
  - Equi-depth/height**: equal count ("mass") per bucket
  - V-Optimal**: minimize total variance of value counts in buckets

10

## Answering Queries Using Histograms

- Answering queries from 1-D histograms (in general):
  - (Implicitly) map the histogram back to an approximate relation, & apply the query to the approximate relation
- Inside each bucket: **Uniformity Assumption**
  - Continuous value mapping
  - Uniform spread mapping

11

## Haar Wavelet Synopses

- Wavelets**: mathematical tool for hierarchical decomposition of functions/signals
- Haar wavelets**: simplest wavelet basis, easy to understand and implement
  - Recursive pairwise averaging and differencing** at different resolutions

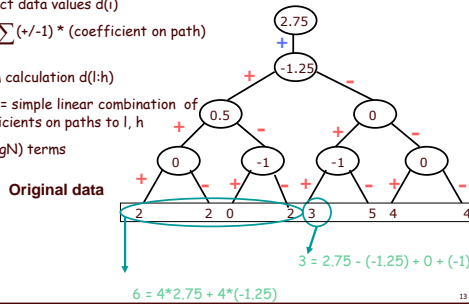
Resolution	Averages	Detail Coefficients
3	$D = [2, 2, 0, 2, 3, 5, 4, 4]$	----
2	$[2, 1, 4, 4]$	$[0, -1, -1, 0]$
1	$[1.5, 4]$	$[0.5, 0]$
0	$[2.75]$	$[-1.25]$

**Haar wavelet decomposition:**  $[2.75, -1.25, 0.5, 0, 0, -1, -1, 0]$

12

## Haar Wavelet Coefficients

- Hierarchical decomposition structure (a.k.a. **Error Tree**)
  - Conceptual tool to "visualize" coefficient supports & data reconstruction
- Reconstruct data values  $d(i)$ 
  - $d(i) = \sum (+/-1) * (\text{coefficient on path})$
- Range sum calculation  $d(l:h)$ 
  - $d(l:h)$  = simple linear combination of coefficients on paths to  $l, h$
- Only  $O(\log N)$  terms




## Wavelet Data Synopses

- Compute Haar wavelet decomposition of  $D$
- **Coefficient thresholding**: only  $B \ll |D|$  coefficients can be kept
  - $B$  is determined by the available synopsis space
- Approximate query engine can do all its processing over such compact coefficient synopses (joins, aggregates, selections, etc.)
- **Conventional thresholding**: Take  $B$  largest coefficients in *absolute normalized value*
  - Normalized Haar basis: divide coefficients at resolution  $j$  by  $\sqrt{2^j}$
  - All other coefficients are ignored (assumed to be zero)
  - *Provably optimal* in terms of the overall Sum-Squared (L2) Error


## Multi-dimensional Data Synopses

- **Problem**: Approximate the *joint data distribution* of multiple attributes
- **Motivation**
  - Selectivity estimation for queries with multiple predicates
  - Approximating general relations
- **Conventional approach**: Attribute-Value Independence (AVI) assumption
  - $sel(p(A1) \& p(A2) \& \dots) = sel(p(A1)) * sel(p(A2)) * \dots$
  - Simple -- one-dimensional marginals suffice
  - **BUT**: almost always inaccurate, gross errors in practice



## Multi-dimensional Histograms

- Use small number of multi-dimensional buckets to *directly* approximate the joint data distribution
- Uniform spread & frequency approximation within buckets
  - $n(i)$  = no. of distinct values along  $A_i$ ,  $F$  = total bucket frequency
  - approximate data points on a  $n(1)*n(2)*\dots$  uniform grid, each with frequency  $F / (n(1)*n(2)*\dots)$

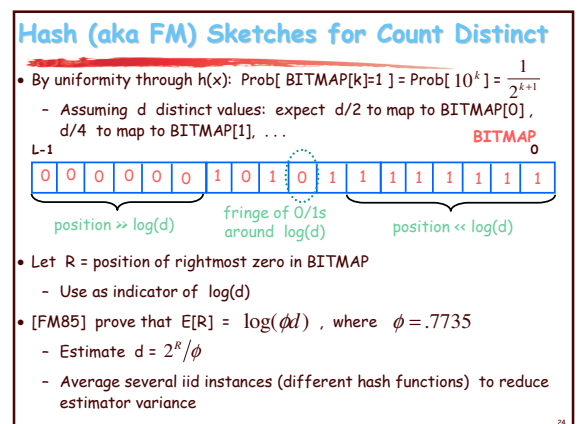
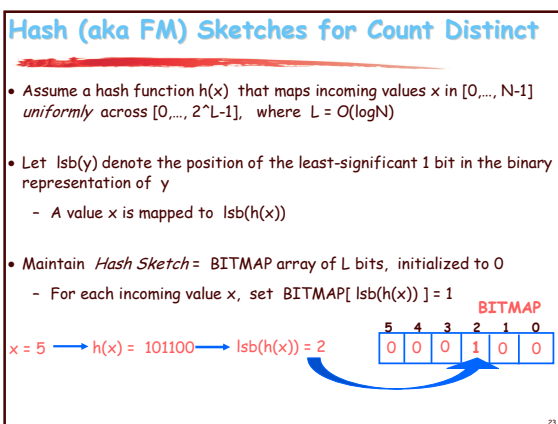
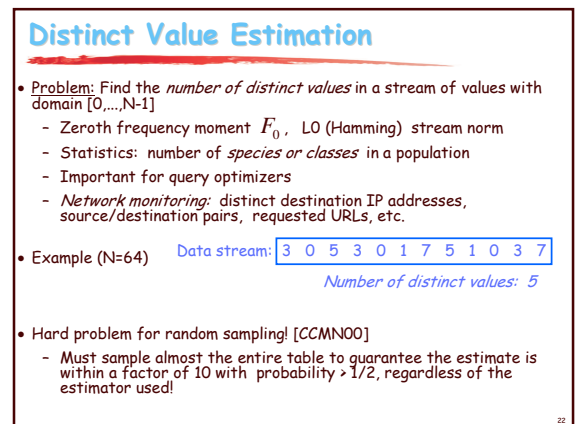
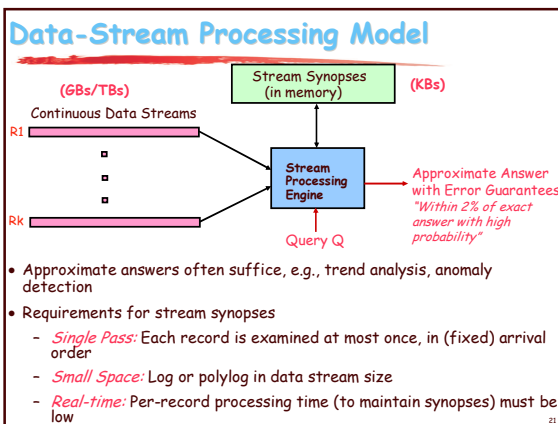
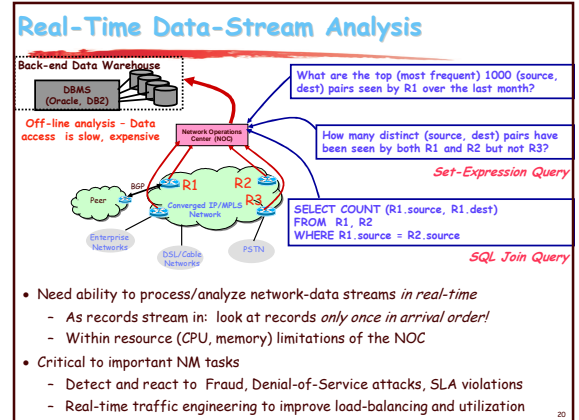
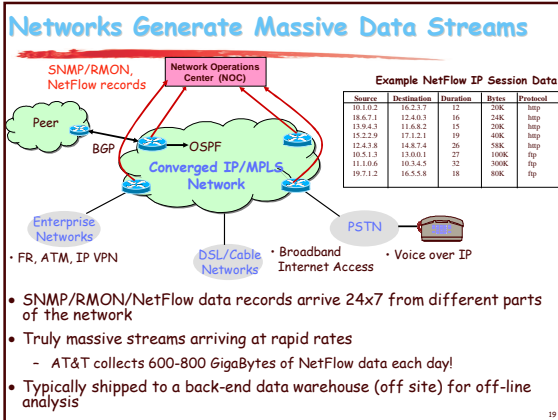


## Data Synopses in Commercial DBMSs

- Sampling operators and 1-D histograms are available in most commercial DBMSs
  - Oracle, DB2, SQL Server, ...
  - Used internally but also exposed to user (e.g., store "sample view")
  - SQL Server has support for 2-D histograms!
- The next step: *Synopses for XML!*
  - How do you effectively summarize a graph structure for queries like `"/a/b[d]/*c" ??`

## Data-Stream Management

- **Traditional DBMS** - data stored in *finite, persistent data sets*
- **Data Streams** - distributed, continuous, unbounded, rapid, time varying, noisy, ...
- **Data-Stream Management** - variety of modern applications
  - Network monitoring and traffic engineering
  - Telecom call-detail records
  - Network security
  - Financial applications
  - Sensor networks
  - Web logs and clickstreams



### A Little Streaming Puzzle...

- *Input*: A stream of  $N$  numbers/elements
- *Output*: The stream *majority element* (if one exists)
  - $e$  is a majority element if  $\text{frequency}(e) > N/2$
- Q: How do you do this in *small space*??
  - Hint: Use just *two* memory locations
  - Hint++: Look at this as a "knockout tournament"
- Feeling adventurous?
  - How do you do the same majority query over a stream of *insertions and deletions*?
  - *Input*: Stream of  $\langle e, + \rangle$  = insert  $e$  ,  $\langle e, - \rangle$  = delete  $e$
  - Hint: Use a little more memory...

25

### In Summary: Not your parents' DBMS!

- Database/data-management research goes far beyond the basics!
- Extends from distributed systems to theory to approximation algorithms to probability/statistics to ...
  - Applications: data mining, sensornets, p2p, ...
  - Just pick up a copy of recent SIGMOD/VLDB proceedings
- More and more relevant in dealing with the "*data tsunami*"
  - Data is everywhere! And, it's constantly growing in volume!
- Exciting, relevant research!

26

### More details...

- Tutorial slides on approximate query processing and data streams

<http://www2.berkeley.intel-research.net/~minos/tutorials.html>

27