Processing Complex Aggregate Queries over Data Streams

SIGMOD 2002

Alin Dobra
Minos Garofalakis
Johannes Gehrke
Rajeev Rastogi

June 4, 2002

Processing Network Data Streams

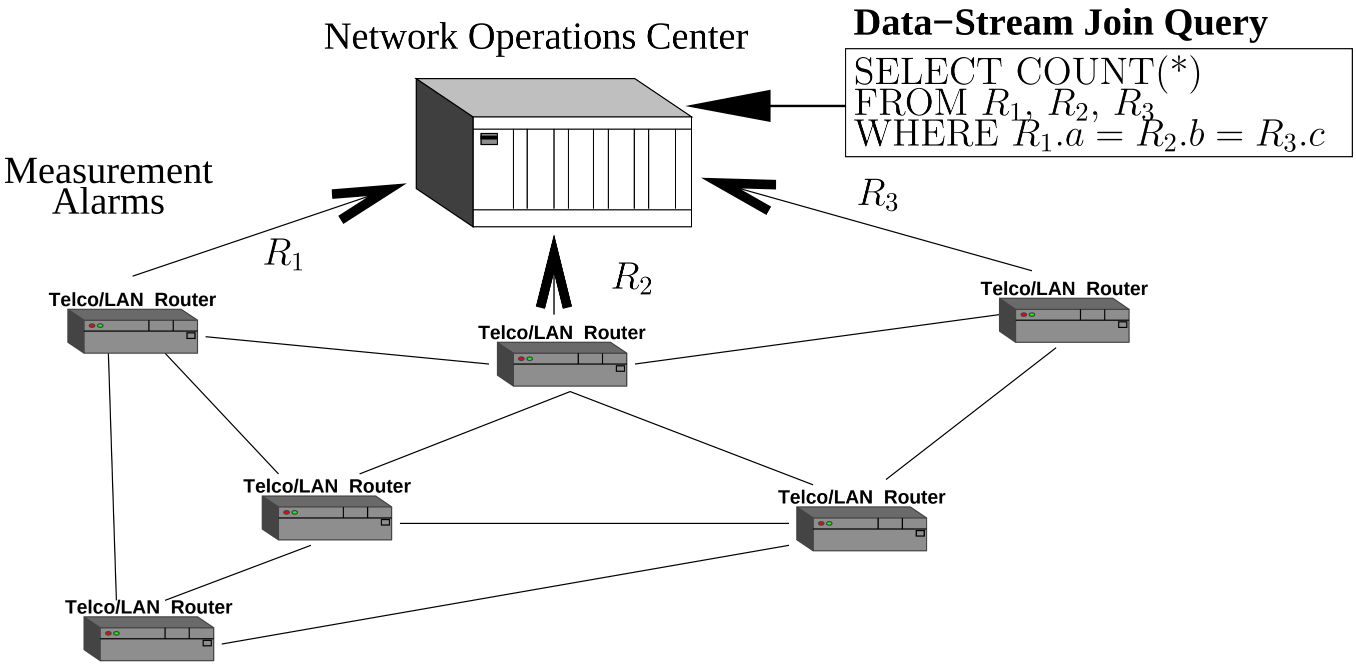

Computations over Streaming Data

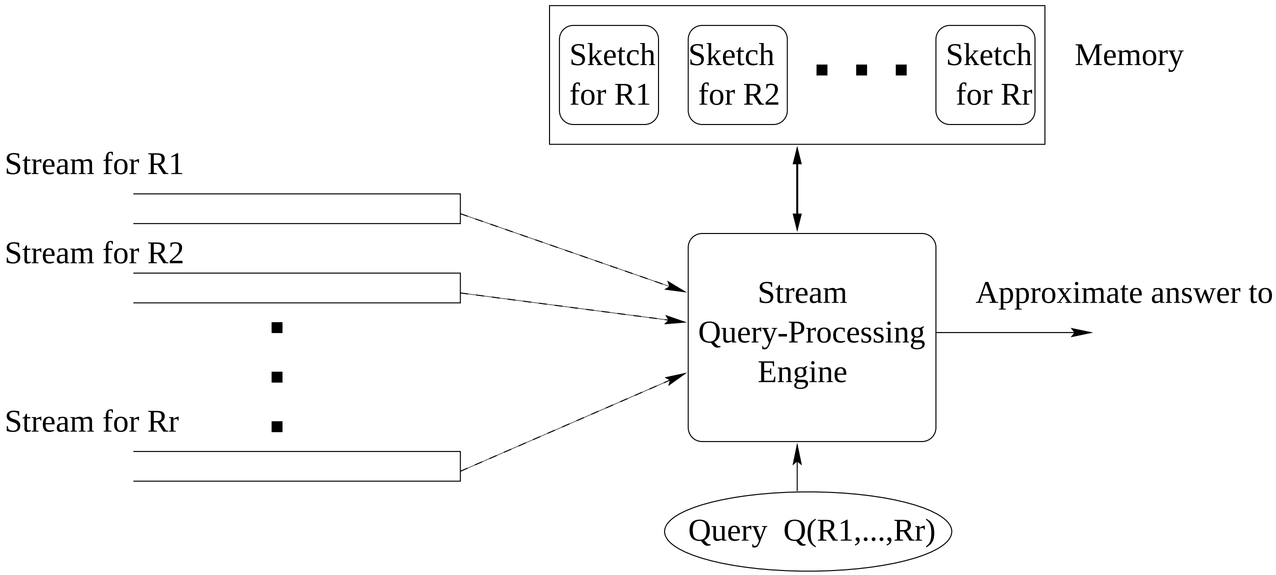

- **Goal:** Approximately answer JOIN-COUNT and JOIN-SUM queries over streams

Outline of the Talk

- Motivation
- Sketch-based randomized algorithms
- Sketch-based approximation of aggregate queries results
- Sketch-partitioning for estimation accuracy boosting
- Experimental evaluation
- Summary

Sketch-Based Randomized Algorithms [AMS96]

- Estimate $F(\mathcal{D})$ for some function F and some data $\mathcal{D}$

Method:

- Build a **probability space** and a **random variable** X with the properties:

  1) $E[X] = F(\mathcal{D}) \geq L_E$

  2) $\text{Var}(X) \leq U_V$

- Combine samples of X to achieve relative error ϵ with probability at least $1 - \delta$
- Boost accuracy to ϵ by averaging $\frac{8U_V}{\epsilon^2 L_E^2}$ pairwise independent samples of X
- Boost confidence to $1 - \delta$ by taking the median of $2\log(1/\delta)$ averages

Example usage: frequency moments [AMS96], size of join [AGMS99], L^1 norm [FKSV99], wavelet decomposition [GKMS01]

Sketch-Based Randomized Algorithms (cont.)

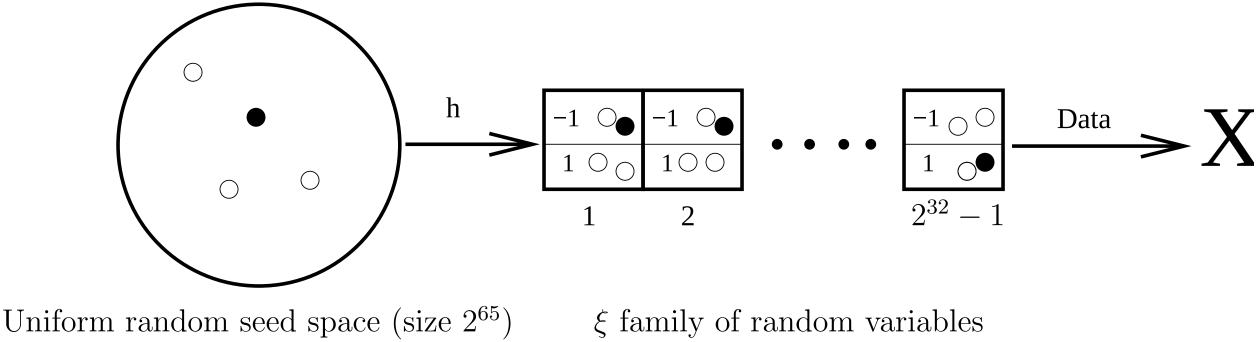

- $\xi_i(s) = h(s,i) \in \{-1,+1\}$
- family ξ is 4-wise independent, i.e.

$$\forall i_1 \neq i_2 \neq i_3 \neq i_4, \forall v_1, v_2, v_3, v_4 \in \{-1,+1\},$$
$$P[\xi_{i_1} = v_1 \wedge \xi_{i_2} = v_2 \wedge \xi_{i_3} = v_3 \wedge \xi_{i_4} = v_4] =$$
$$P[\xi_{i_1} = v_1] P[\xi_{i_2} = v_2] P[\xi_{i_3} = v_3] P[\xi_{i_4} = v_4]$$

Estimation of COUNT$(F \bowtie_a G)$ [AGMS99]

F

$\cdots$	a	$\cdots$
	1	
	1	
	2	
	3	
	1	
	3	

$\Rightarrow$

i	f_i
1	3
2	1
3	2

G

$\cdots$	a	$\cdots$
	3	
	3	
	1	
	1	
	1	

$\Rightarrow$

i	g_i
1	3
2	0
3	2

- Estimate COUNT$(F \bowtie_a G) = \sum_{i=1}^{3} f_i g_i = 3 \cdot 3 + 1 \cdot 0 + 2 \cdot 2 = 13$

Estimation of COUNT$(F \bowtie_a G)$ (cont.)

i	1	2	3
ξ_i	-1	$+1$	-1

F_a	ξ_a	$X_F = \sum_{t \in F} \xi_{t.a}$
1	-1	-1
1	-1	-2
2	$+1$	-1
3	-1	$+0$
1	-1	-1
3	-1	-2

G_a	ξ_a	$X_G = \sum_{t \in G} \xi_{t.a}$
3	-1	-1
3	-1	-2
1	-1	-3
1	-1	-4
1	-1	-5

$$X = X_F X_G = -2 \cdot -5 = 10 \approx 13$$

$$\mathsf{SJ}(F) = (3 \cdot 3) + (1 \cdot 1) + (2 \cdot 2) = 14, \quad \mathsf{SJ}(G) = 13$$

Estimation of COUNT$(F \bowtie_a G)$ (cont.)

- To estimate COUNT$(F \bowtie G) = \sum_{i=1}^{n} f_i g_i$ define:

$$X_F = \sum_{i=1}^{n} f_i \xi_i = \sum_{t \in F} \xi_{t.a}$$

$$X_G = \sum_{i=1}^{n} g_i \xi_i = \sum_{t \in G} \xi_{t.a}$$

- With $X = X_F X_G$ we have:

$$E[X] = E\Big[\sum_{i=1}^{n} f_i g_i \xi_i^2 + \sum_{i \neq i'} f_i g_{i'} \xi_i \xi_{i'}\Big]$$

$$= \mathsf{COUNT}(F \bowtie_a G)$$

$$\mathsf{Var}(X) \leq 2\ \mathsf{SJ}(F)\ \mathsf{SJ}(G)$$

Outline of the Talk

- Motivation
- Sketch-based randomized algorithms
- Sketch-based approximation of aggregate queries results
- Sketch-partitioning for estimation accuracy boosting
- Experimental evaluation
- Summary

Using Sketches to Answer SUM Queries

- Estimate SUM$_b\left(F(a) \bowtie_a G(a,b)\right) = \sum_{i=1}^{3} f_i \left(\sum_{t\in g, t.a=i} t.b\right)$

i	1	2	3
ξ_i	-1	$+1$	-1

F_a	ξ_a	$X_F = \sum_{t\in F} \xi_{t.a}$
1	-1	-1
1	-1	-2
2	$+1$	-1
3	-1	$+0$
1	-1	-1
3	-1	-2

G_a	G_b	ξ_a	$X_G = \sum_{t\in G} t.b\ \xi_{t.a}$
3	2	-1	-2
3	2	-1	-4
1	1	-1	-5
1	2	-1	-7
1	1	-1	-8

$$X = X_F X_G = -2 \cdot -8 = 16 \approx 20$$

Using Sketches to Answer SUM Queries (cont.)

- To estimate $\mathsf{SUM}_b\left(F(a) \bowtie_a G(a,b)\right) = \sum_{i=1}^{n} f_i \left(\sum_{t \in g, t.a=i} t.b\right)$ define:

$$X_F = \sum_{i=1}^{n} f_i \xi_i = \sum_{t \in F} \xi_{t.a}$$

$$X_G = \sum_{i=1}^{n} \Big(\sum_{t \in G, t.a=i} t.b \Big) \xi_i = \sum_{t \in G} t.b \ \xi_{t.a}$$

- With $X = X_F X_G$

$$E[X] = \mathsf{SUM}_b\left(F(a) \bowtie_a G(a,b)\right)$$

$$\mathsf{Var}(X) \leq 2\ \mathsf{SJ}(F) \sum_{i=1}^{n} \Big(\sum_{t \in G, t.a=i} t.b \Big)^2$$

Extension to COUNT$(F \bowtie_a G \bowtie_b H)$

- Key idea: use independent ξ families for each join attribute

i	1	2	3
ξ_i^a	-1	$+1$	-1

j	1	2
ξ_j^b	$+1$	-1

F_a	$\xi_{t.a}^a$	$X_F = \sum \xi_{t.a}^a$
1	-1	-1
1	-1	-2
2	$+1$	-1
3	-1	$+0$
1	-1	-1
3	-1	-2

G_a	G_b	$\xi_{t.a}^a$	$\xi_{t.b}^b$	$X_G = \sum \xi_{t.a}^a \xi_{t.b}^b$
3	2	-1	-1	-1
3	2	-1	-1	-2
1	1	-1	$+1$	-1
1	2	-1	-1	0
1	1	-1	$+1$	1

H_b	$\xi_{t.b}^b$	$X_H = \sum \xi_{t.b}^b$
2	-1	-1
2	-1	-2
1	$+1$	-1
2	-1	-2

$$X = X_F X_G X_H = -2 \cdot 1 \cdot -2 = 4 \approx 21$$

Extention to COUNT$(F \bowtie_a G \bowtie_b H)$(cont.)

- To estimate

$$\text{COUNT}(F \bowtie_a G \bowtie_b H) = \sum_{i=1}^{n_1} \sum_{j=1}^{n_2} f_i g_{ij} h_j$$

- Define:

$$X_F = \sum_{i=1}^{n_1} f_i \xi_i^a, \quad X_G = \sum_{i=1}^{n_1} \sum_{j=1}^{n_2} g_{ij} \xi_i^a \xi_j^b, \quad X_H = \sum_{j=1}^{n_2} h_j \xi_j^b$$

- If ξ^a and ξ^b are independent families of ± 1 4-wise independent pseudo random variables

$$E[X_F X_G X_H] = \text{COUNT}(F \bowtie_a G \bowtie_b H)$$
$$\text{Var}(X_F X_G X_H) \leq 4 \ \text{SJ}(F) \ \text{SJ}(G) \ \text{SJ}(H)$$

Estimation of COUNT$(R_1 \bowtie \cdots \bowtie R_r)$

- For each of the n equality join constraint build independent family of pseudo random variables
- For every relation $R_l(a_1, \ldots, a_m)$ compute samples of the random variable X_{R_l} defined as:

$$X_{R_l} = \sum_{i_1}^{n_1} \cdots \sum_{i_m}^{n_m} f_{i_1,\ldots,i_m} \xi_{1,i_1} \ldots \xi_{m,i_m} = \sum_{t \in R} \xi_{1,t.a_1} \cdots \xi_{m,t.a_m}$$

$$X = \prod_{l=1}^{r} X_{R_l}$$

- Can show:

$$E[X] = \mathsf{COUNT}(R_1 \bowtie \cdots \bowtie R_r)$$

$$\mathsf{Var}(X) \leq 2^{2n} \prod_{l=1}^{r} \mathsf{SJ}(R_l)$$

Outline of the Talk

- Motivation
- Sketch-based randomized algorithms
- Sketch-based approximation of aggregate queries results
- Sketch-partitioning for estimation accuracy boosting
- Experimental evaluation
- Summary

Sketch Partitioning

Problem: large variance $\Rightarrow$ loose estimation guarantees.

Our solution: sketch partitioning

i	f_i	g_i
1	20	2
2	5	15
3	10	3
4	2	10

$$\begin{aligned}\mathsf{Var}(X) &\approx 2\ \mathsf{SJ}(F)\ \mathsf{SJ}(G)\\ &= 2(20^2+5^2+10^2+2^2)(2^2+15^2+3^2+10^2)\\ &= 357604\end{aligned}$$

Idea: split domain $I = \{1,2,3,4\}$ into $I_1 = \{1,3\}$ and $I_2 = \{2,4\}$

- F splits into F_1 and F_2, G into G_1 and G_2
- build X_1 to estimate COUNT$(F_1 \bowtie G_1)$ and independently X_2 to estimate COUNT$(F_2 \bowtie G_2)$
- take $X' = X_1 + X_2$; have $E[X'] =$ COUNT$(F \bowtie G)$

Sketch Partitioning (cont.)

- Estimation of COUNT$(F_1 \bowtie G_1)$

i	f_i	g_i
1	20	2
3	10	3

$$\begin{aligned} \mathsf{Var}(X_1) &\approx 2\ \mathsf{SJ}(F_1)\ \mathsf{SJ}(G_1) \\ &= 2(20^2+10^2)(2^2+3^2) \\ &= 13000 \end{aligned}$$

- Estimation of COUNT$(F_2 \bowtie G_2)$

i	f_i	g_i
2	5	15
4	2	10

$$\begin{aligned} \mathsf{Var}(X_2) &\approx 2\ \mathsf{SJ}(F_2)\ \mathsf{SJ}(G_2) \\ &= 2(5^2+2^2)(15^2+10^2) \\ &= 18850 \end{aligned}$$

- $\mathsf{Var}(X') = \mathsf{Var}(X_1) + \mathsf{Var}(X_2) = 31850$
- Improvement

$$\frac{\mathsf{Var}(X)/2}{\mathsf{Var}(X')} = \frac{357604/2}{31850} \approx 5.6$$

Binary Sketch Partitioning

- Prior information: historical data, histograms.
- Find the partitioning $I = I_1 \cup I_2$ and the space allocation $m = m_1 + m_2$ that minimizes
$$\frac{\mathsf{Var}(X_1)}{m_1} + \frac{\mathsf{Var}(X_2)}{m_2},$$
where
$$\mathsf{Var}(X_k) \approx 2 \sum_{i \in I_k} f_i^2 \sum_{i \in I_k} g_i^2.$$
- Allocate space proportional to $\sqrt{\mathsf{Var}(X_k)}$. In example 5:6
- Have to look only at partitioning in the order f_i/g_i to find optimum $\Rightarrow O(|\,I\,|)$
- In example order is $\{1,3,2,4\}$. Optimal partition is $\{1,3\} \cup \{2,4\}$.

K-ary Sketch Partitioning

- Want to split domain of join attribute in K parts
- Allocate space proportional to $\sqrt{\text{Var}(X_k)}$
- Have to look only at partitioning in the order f_i/g_i to find optimum (generalization of previous result)
- Dynamic programming gives solution in time $O(K\,|\,I\,|^2)$ and space $O(K\,|\,I\,|)$
- Approximate frequencies with histograms
  - time and space dependency on number of buckets instead of $|\,I\,|$
  - provable approximation quality
- Generalization to larger joins possible: *details in the paper*

Experimental Study

Datasets:

- Census data set (`www.bls.census.gov`):
  - Current Population Survey data for Aug 1999(72100) and Aug 2001(81600)
  - Attributes used:
    - income(1:14)
    - education(1:46)
    - age(1:99)
    - weekly_wage and weekly_ wage_overtime(0:288416)

Comparison: estimation using unidimensional equi-depth histograms

Query load: JOIN-COUNT queries relations

Error metric: relative error $= 100 \frac{|\text{actual} - \text{approx}|}{\text{actual}} \%$

Sketches v/s Histograms: Census data

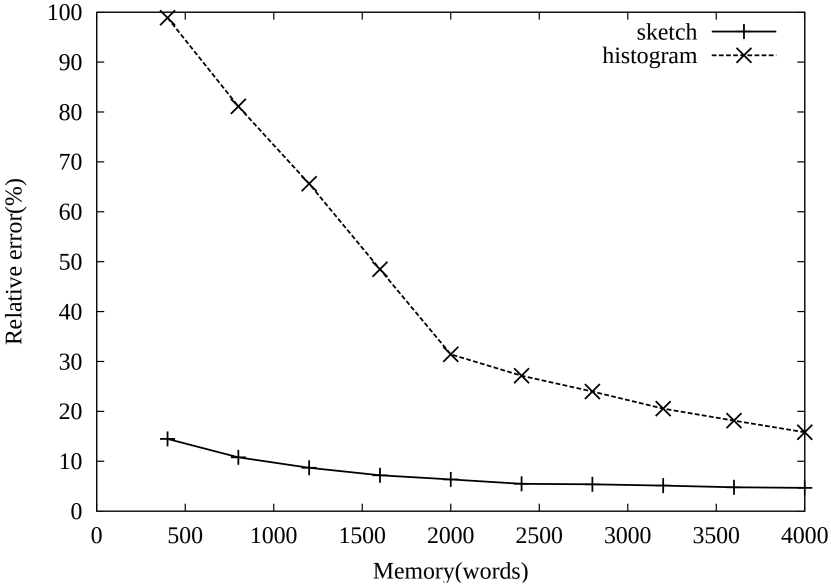

Census1999.weekly_wage = Census2001.weekly_wage

Sketches v/s Histograms: Census data (cont.)

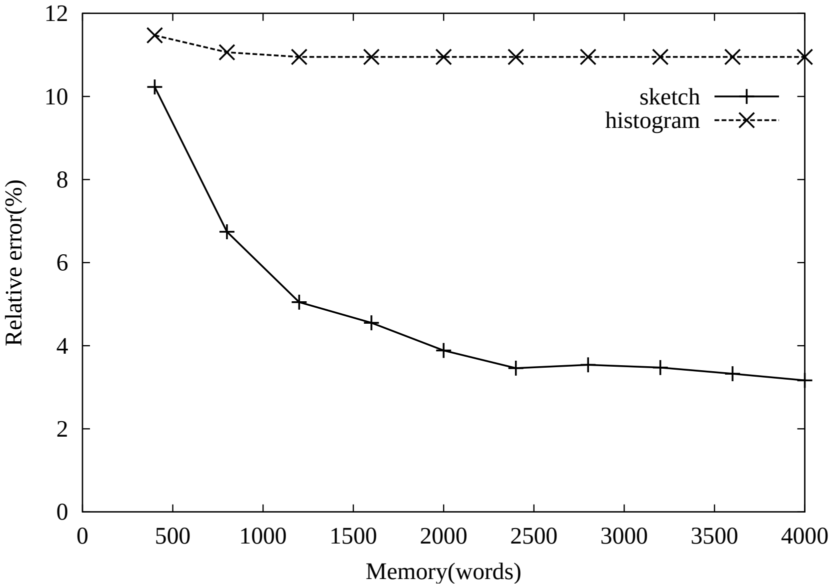

Census1999.age = Census2001.age

Census1999.education = Census2001.education

Sketches v/s Histograms: Census data (cont.)

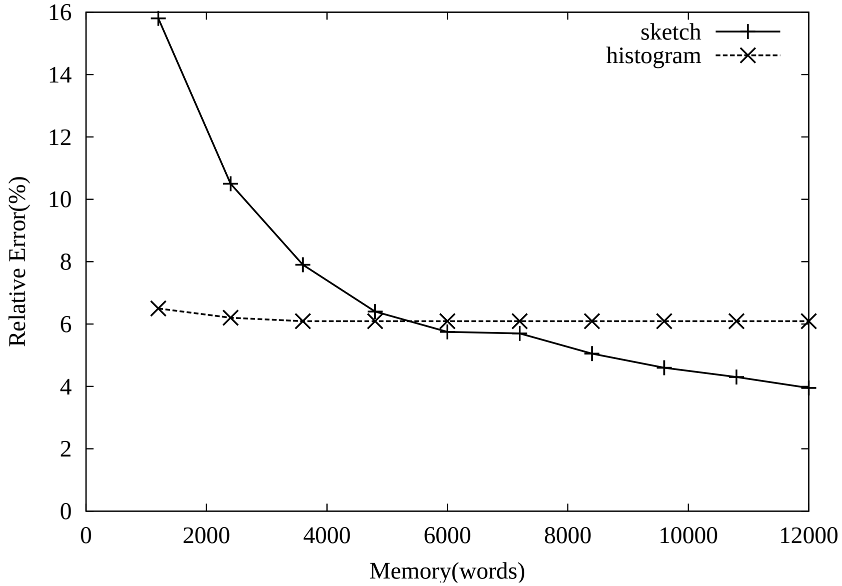

Star join of four copies of Census 2001 on age, education and income

Sketch Partitioning: Census Data Sets

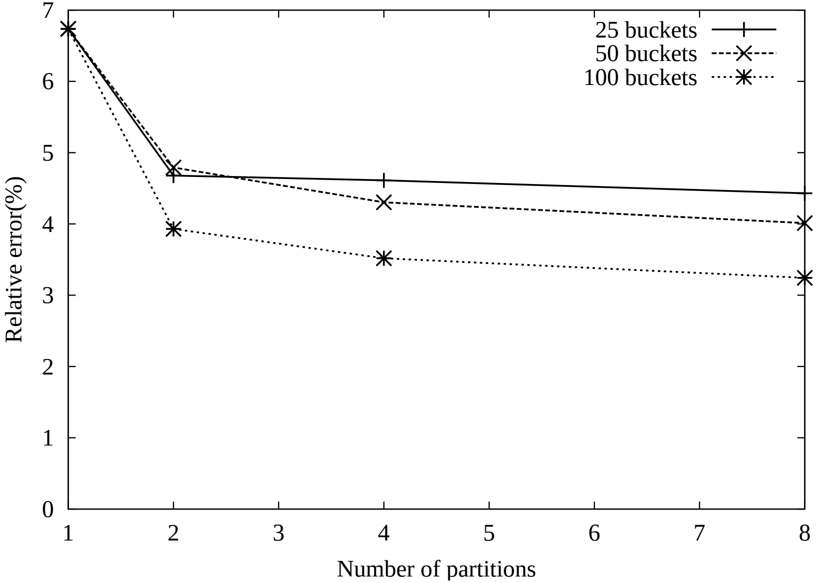

Census1999.weekly_wage_overtime = Census2001.weekly_wage_overtime

Sketch Partitioning: Census Data Sets (cont.)

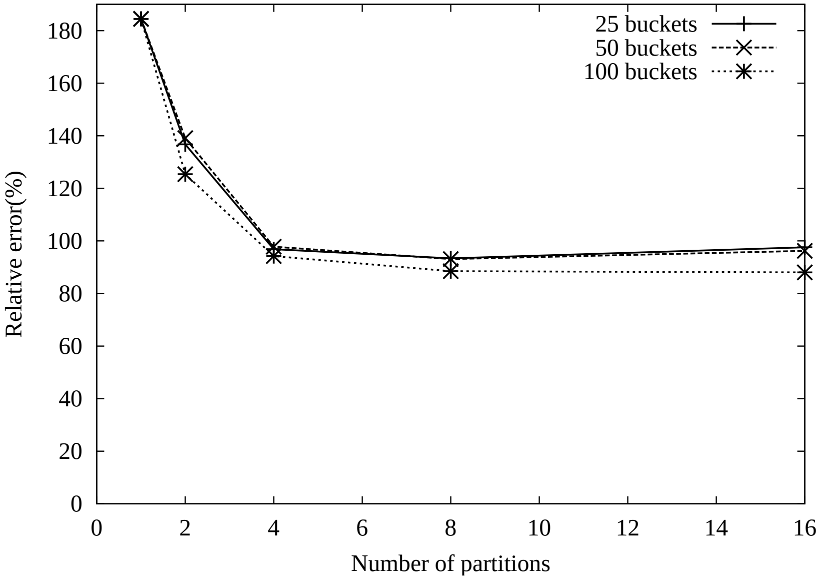

Census1999.weekly_wage_overtime = Census2001.weekly_wage

Census1999.weekly_wage = Census2001.weekly_wage_overtime

Summary

- Shown how to process multi-join decision support queries over streams
- Proposed sketch partitioning – improves estimate guarantees
- Shown experimental evidence that the proposed techniques work in practice