


# Computing Statistical Summaries over Massive Distributed Data

Ke (Kevin) Yi

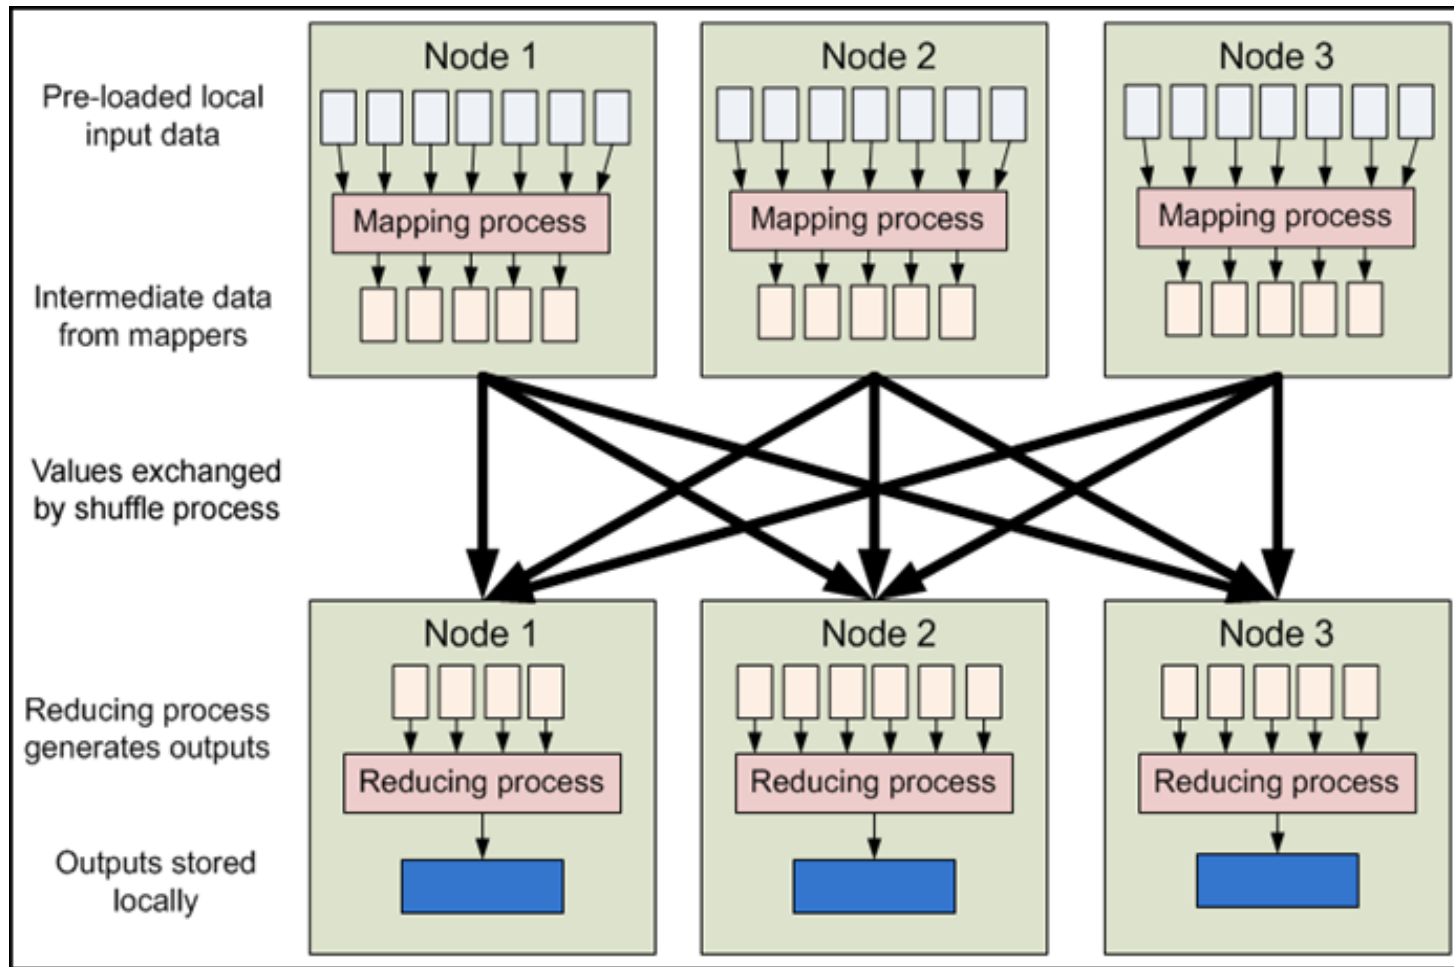
Hong Kong University of Science and Technology

# Distributed Systems for Massive Data: MapReduce



Open source implementation: Hadoop


# Distributed Systems for Massive Data: MapReduce



Open source implementation: Hadoop


Suitable for batch processing (e.g., index construction)

# Distributed Systems for Massive Data: Dremel



No open source implementation yet


# Distributed Systems for Massive Data: Dremel




No open source implementation yet

Suitable for analytical queries (e.g., extracting a summary)

# (Simplified) Model of Computation




# (Simplified) Model of Computation



- The root broadcasts a message to initialize computation
- Each node computes a summary on its local data
- The root combines the summaries to produce a global summary

# (Simplified) Model of Computation




- The root broadcasts a message to initialize computation
- Each node computes a summary on its local data
- The root combines the summaries to produce a global summary
- Using minimum communication (and load balancing)



# Outline

- Model of computation
- Frequency estimation (heavy hitters)
- Quantiles (order statistics)
- Other problems

# Problem: Frequency Estimation



- Input: Multiset  $S$  of  $N$  items drawn from the universe  $[u] = \{1 \dots u\}$   
For example, all IP addresses
- Each node  $j \in [k]$  holds a subset of  $S$   
For any item  $i \in [u]$   
 $x_{ij}$ : total number of  $i$ 's in node  $j$  (local count)  
 $y_i = \sum_{j=1}^k x_{ij}$  (global count)
- Compute  $y_i$  for each  $i$

# Frequency Estimation: Possible Solutions

- Compute exactly: **send everything**

# Frequency Estimation: Possible Solutions

- Compute exactly: **send everything**
- Approximate each  $y_i$  within additive error  $\epsilon N$

# Frequency Estimation: Possible Solutions

- Compute exactly: **send everything**
- Approximate each  $y_i$  within additive error  $\epsilon N$
- Sketching: Each node computes a sketch of its own data and sends it to the coordinator.

*Count-min sketch, MG sketch, Space saving, etc.*

Sketch size:  $O(1/\epsilon)$

Communication cost:  $O(k/\epsilon)$

# Frequency Estimation: Possible Solutions

- Compute exactly: **send everything**
- Approximate each  $y_i$  within additive error  $\epsilon N$
- Sketching: Each node computes a sketch of its own data and sends it to the coordinator.

*Count-min sketch, MG sketch, Space saving, etc.*

Sketch size:  $O(1/\epsilon)$

Communication cost:  $O(k/\epsilon)$

- Random sampling

Uniformly randomly sample a subset of size  $O(1/\epsilon^2)$

# Frequency Estimation: Possible Solutions

- Compute exactly: **send everything**
- Approximate each  $y_i$  within additive error  $\epsilon N$
- Sketching: Each node computes a sketch of its own data and sends it to the coordinator.

*Count-min sketch, MG sketch, Space saving, etc.*

Sketch size:  $O(1/\epsilon)$

Communication cost:  $O(k/\epsilon)$

- Random sampling

Uniformly randomly sample a subset of size  $O(1/\epsilon^2)$

- We can achieve:  $O(\sqrt{k}/\epsilon)$

Typical values of  $\epsilon = 10^{-3} \sim 10^{-6}$ ,  $k = 10^2 \sim 10^4$

We assume  $k < 1/\epsilon^2$

# HT estimator [Horvitz and Thompson 56]

Each node holds a set of (item, count) pairs

(item, count)

(1, 20)  
(2, 13)  
(3, 35)  
(4, 12)  
(5, 5)  
(6, 22)

node  $j$


send each pair  $(i, x_{ij})$  with probability  $g(x_{ij})$



# HT estimator [Horvitz and Thompson 56]


Each node holds a set of (item, count) pairs

(item, count)



node  $j$

send each pair  $(i, x_{ij})$  with probability  $g(x_{ij})$



# HT estimator [Horvitz and Thompson 56]

HT estimator for  $x_{ij}$ :

$$Y_{i,j} = \frac{x_{i,j}}{g(x_{i,j})} \text{ if it is sampled, otherwise } 0$$

This is an unbiased estimator

Estimator for  $y_i$ :

$$Y_i = Y_{i,1} + \cdots + Y_{i,n}$$

# HT estimator [Horvitz and Thompson 56]

HT estimator for  $x_{ij}$ :

$$Y_{i,j} = \frac{x_{i,j}}{g(x_{i,j})} \text{ if it is sampled, otherwise } 0$$

This is an unbiased estimator

Estimator for  $y_i$ :

$$Y_i = Y_{i,1} + \cdots + Y_{i,n}$$

$$\begin{aligned} \text{Var}[Y_{i,j}] &= \left( \frac{x_{i,j}}{g(x_{i,j})} - x_{i,j} \right)^2 g(x_{i,j}) + (x_{i,j})^2 (1 - g(x_{i,j})) \\ &= \frac{x_{i,j}^2 (1 - g(x_{i,j}))}{g(x_{i,j})} \end{aligned}$$

$$\text{Var}[Y_i] = \sum_{j=1}^n \text{Var}[Y_{ij}] = \sum_{j=1}^n \frac{x_{i,j}^2 (1 - g(x_{i,j}))}{g(x_{i,j})}$$

# Sampling Function

**Question:** What sampling function  $g(x)$  should we use

# Sampling Function

**Question:** What sampling function  $g(x)$  should we use

Accuracy: standard deviation less than  $\epsilon N$

A function is **valid**, if  $\text{Var}[Y_i] \leq (\epsilon N)^2$  for all items  $i$

# Sampling Function

**Question:** What sampling function  $g(x)$  should we use

Accuracy: standard deviation less than  $\epsilon N$

A function is **valid**, if  $\text{Var}[Y_i] \leq (\epsilon N)^2$  for all items  $i$

Communication cost:  $\sum_{i,j} g(x_{ij})$

# Sampling Function

**Question:** What sampling function  $g(x)$  should we use

Accuracy: standard deviation less than  $\epsilon N$

A function is **valid**, if  $\text{Var}[Y_i] \leq (\epsilon N)^2$  for all items  $i$

Communication cost:  $\sum_{i,j} g(x_{ij})$

Optimal valid  $g(x)$ ?

# A Worst-Case Optimal Sampling Function

$$g_1(x) = \min\left\{\frac{\sqrt{k}}{\varepsilon N}x, 1\right\}$$



# A Worst-Case Optimal Sampling Function

$$g_1(x) = \min\left\{\frac{\sqrt{k}}{\varepsilon N}x, 1\right\}$$

Can show:

$$\blacksquare \quad \text{Var}[Y_i] = - \left( \frac{y_i}{\sqrt{k}} - \frac{\varepsilon N}{2} \right)^2 + \frac{(\varepsilon N)^2}{4} \leq \frac{1}{4}(\varepsilon N)^2,$$

i.e.,  $g_1(x)$  is valid

- Communication cost of using  $g_1(x)$  is  $O(\sqrt{k}/\varepsilon)$
- Communication cost of any valid sampling function is  $\Omega(\sqrt{k}/\varepsilon)$  in the worst case (i.e., on some input)

# A Worst-Case Optimal Sampling Function

$$g_1(x) = \min\left\{\frac{\sqrt{k}}{\varepsilon N}x, 1\right\}$$

Can show:

$$\blacksquare \quad \text{Var}[Y_i] = - \left( \frac{y_i}{\sqrt{k}} - \frac{\varepsilon N}{2} \right)^2 + \frac{(\varepsilon N)^2}{4} \leq \frac{1}{4}(\varepsilon N)^2,$$

i.e.,  $g_1(x)$  is valid

- Communication cost of using  $g_1(x)$  is  $O(\sqrt{k}/\varepsilon)$
- Communication cost of any valid sampling function is  $\Omega(\sqrt{k}/\varepsilon)$  in the worst case (i.e., on some input)
- A very recent result shows that **any** algorithm has to spend  $\Omega(\sqrt{k}/\varepsilon)$  bits of communication in the worst case [Woodruff, Zhang, manuscript]

# Another Sampling Function

$$g_2(x) = (g_1(x))^2$$

# Another Sampling Function

$$g_2(x) = (g_1(x))^2$$

Can show:

- $g_2$  is also valid

# Another Sampling Function

$$g_2(x) = (g_1(x))^2$$

Can show:

- $g_2$  is also valid
- Clearly,  $g_1(x) \geq g_2(x)$

$g_1$ 's communication cost is always  $\Theta(\sqrt{k}/\varepsilon)$ , while  $g_2$  can be much better when there are many small local counts

# Another Sampling Function

$$g_2(x) = (g_1(x))^2$$

Can show:

- $g_2$  is also valid
- Clearly,  $g_1(x) \geq g_2(x)$

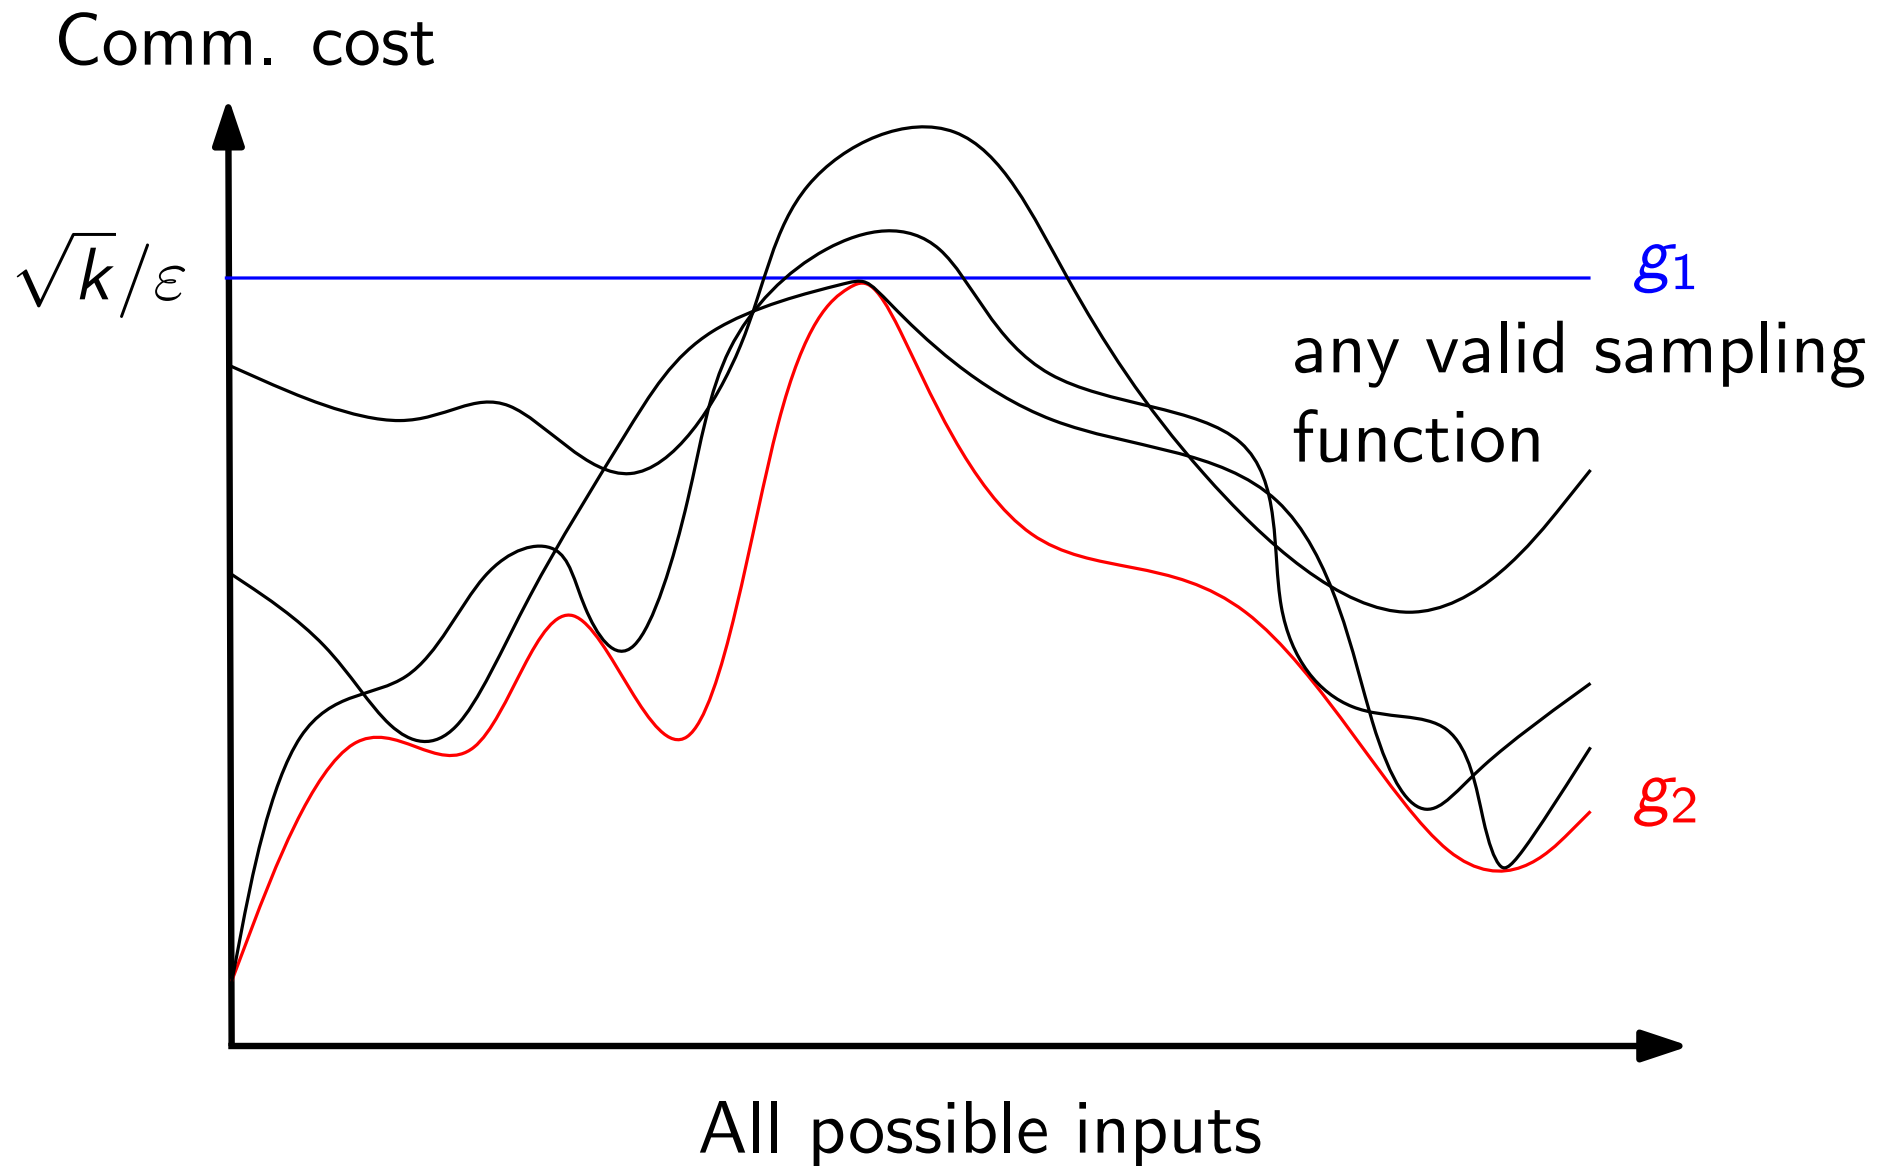
$g_1$ 's communication cost is always  $\Theta(\sqrt{k}/\varepsilon)$ , while  $g_2$  can be much better when there are many small local counts

- A stronger optimality:  $g_2(x)$  is **instance-optimal**

Define  $opt(I) = \sum_{i,j} g_2(x_{i,j})$  on input  $I : \{x_{i,j}\}$

Can show that on every input  $I$ , any valid sampling function must have cost  $\Omega(opt(I))$

# Instance Optimality



# Further Reducing Communication Cost

$$g_1(x) = \min\left\{\frac{\sqrt{k}}{\varepsilon N}x, 1\right\}$$

HT estimator for  $x_{ij}$ :

$$Y_{i,j} = \frac{x_{i,j}}{g(x_{i,j})} \text{ if it is sampled, otherwise } 0$$

Estimator for  $y_i$ :

$$Y_i = Y_{i,1} + \cdots + Y_{i,n}$$



# Further Reducing Communication Cost

$$g_1(x) = \min\left\{\frac{\sqrt{k}}{\varepsilon N}x, 1\right\}$$

HT estimator for  $x_{ij}$ :

$$Y_{i,j} = \frac{x_{i,j}}{g(x_{i,j})} \text{ if it is sampled, otherwise } 0$$

Estimator for  $y_i$ :

$$Y_i = Y_{i,1} + \cdots + Y_{i,n}$$

$$Y_i = \frac{\varepsilon N}{\sqrt{k}}(1 + 0 + 1 + 1 + \cdots + 0 + 1)$$

# Further Reducing Communication Cost

$$g_1(x) = \min\left\{\frac{\sqrt{k}}{\varepsilon N}x, 1\right\}$$

HT estimator for  $x_{ij}$ :

$$Y_{i,j} = \frac{x_{i,j}}{g(x_{i,j})} \text{ if it is sampled, otherwise } 0$$

Estimator for  $y_i$ :

$$Y_i = Y_{i,1} + \cdots + Y_{i,n}$$

$$Y_i = \frac{\varepsilon N}{\sqrt{k}}(1 + 0 + 1 + 1 + \cdots + 0 + 1)$$

Each site  $j$  just needs to tell whether  $i$  is sampled or not!

# Further Reducing Communication Cost

$$g_1(x) = \min\left\{\frac{\sqrt{k}}{\varepsilon N}x, 1\right\}$$

HT estimator for  $x_{ij}$ :

$$Y_{i,j} = \frac{x_{i,j}}{g(x_{i,j})} \text{ if it is sampled, otherwise } 0$$

Estimator for  $y_i$ :

$$Y_i = Y_{i,1} + \cdots + Y_{i,n}$$


$$Y_i = \frac{\varepsilon N}{\sqrt{k}}(1 + 0 + 1 + 1 + \cdots + 0 + 1)$$

Each site  $j$  just needs to tell whether  $i$  is sampled or not!

The set of sampled items can be encoded in a Bloom filter, taking  $O(1)$  bits per item


$\Rightarrow$  total cost =  $O(\sqrt{k}/\varepsilon)$  **bits**

# Bloom Filters



- A Bloom filter needs  $O(\log(1/q))$  bits per item
- No false negatives
- False positive probability =  $q$

# Bloom Filters



- A Bloom filter needs  $O(\log(1/q))$  bits per item
- No false negatives
- False positive probability =  $q$

Change the estimator to

$$Y_i = \frac{\epsilon N}{\sqrt{k}} \cdot \frac{Y_{i,1} + \dots + Y_{i,k} - kq}{1 - q}$$

# Sampling with $g_2(x) = (g_1(x))^2$

- $g_2(x)$  samples  $opt(I)$  (item, count) pairs, which may be much smaller than  $O(\sqrt{k}/\varepsilon)$  on many inputs
- But it is a nonlinear sampling function

Estimator for  $y_i$ :

$$Y_i = \frac{x_{i,1}}{g_2(x_{i,1})} + 0 + 0 + \frac{x_{i,4}}{g_2(x_{i,4})} + \dots + 0 + \frac{x_{i,k}}{g_2(x_{i,k})}$$

# Sampling with $g_2(x) = (g_1(x))^2$

- $g_2(x)$  samples  $opt(I)$  (item, count) pairs, which may be much smaller than  $O(\sqrt{k}/\varepsilon)$  on many inputs
- But it is a nonlinear sampling function

Estimator for  $y_i$ :

$opt(I)$  such terms

$$Y_i = \frac{x_{i,1}}{g_2(x_{i,1})} + 0 + 0 + \frac{x_{i,4}}{g_2(x_{i,4})} + \dots + 0 + \frac{x_{i,k}}{g_2(x_{i,k})}$$

# Sampling with $g_2(x) = (g_1(x))^2$

- $g_2(x)$  samples  $opt(I)$  (item, count) pairs, which may be much smaller than  $O(\sqrt{k}/\epsilon)$  on many inputs
- But it is a nonlinear sampling function

Estimator for  $y_i$ :

$opt(I)$  such terms

$$Y_i = \frac{x_{i,1}}{g_2(x_{i,1})} + 0 + 0 + \frac{x_{i,4}}{g_2(x_{i,4})} + \dots + 0 + \frac{x_{i,k}}{g_2(x_{i,k})}$$

- Use  $g_2(x)$  to perform the sampling locally
- Then use  $g_1(x)$  + Bloom filters to sample the  $\frac{x_{i,j}}{g_2(x_{i,j})}$ 's



# Sampling with $g_2(x) = (g_1(x))^2$

- $g_2(x)$  samples  $opt(I)$  (item, count) pairs, which may be much smaller than  $O(\sqrt{k}/\varepsilon)$  on many inputs
- But it is a nonlinear sampling function

Estimator for  $y_i$ :


$opt(I)$  such terms

$$Y_i = \frac{x_{i,1}}{g_2(x_{i,1})} + 0 + 0 + \frac{x_{i,4}}{g_2(x_{i,4})} + \dots + 0 + \frac{x_{i,k}}{g_2(x_{i,k})}$$

- Use  $g_2(x)$  to perform the sampling locally
- Then use  $g_1(x)$  + Bloom filters to sample the  $\frac{x_{i,j}}{g_2(x_{i,j})}$ 's

Can show this takes  $O\left(opt(I) \log^2\left(\frac{\sqrt{k}}{\varepsilon opt(I)}\right)\right)$  bits

# Simulation Results




$k = 1000$ ,  $N = 10^9$  following Zipf distribution with  $\alpha = 1.2$ .  
Estimate the frequencies of the 100 most popular items. Variance  
computed from 100 runs, and take the worst

# Outline

- Model of computation
- Frequency estimation (heavy hitters)
- **Quantiles (order statistics)**
- Other problems


# Quantiles

In a set of  $n$  values, the  $(r/n)$ -quantile is the value ranked at  $r$ .  
The 0.5-quantile is the **median**.



# Quantiles

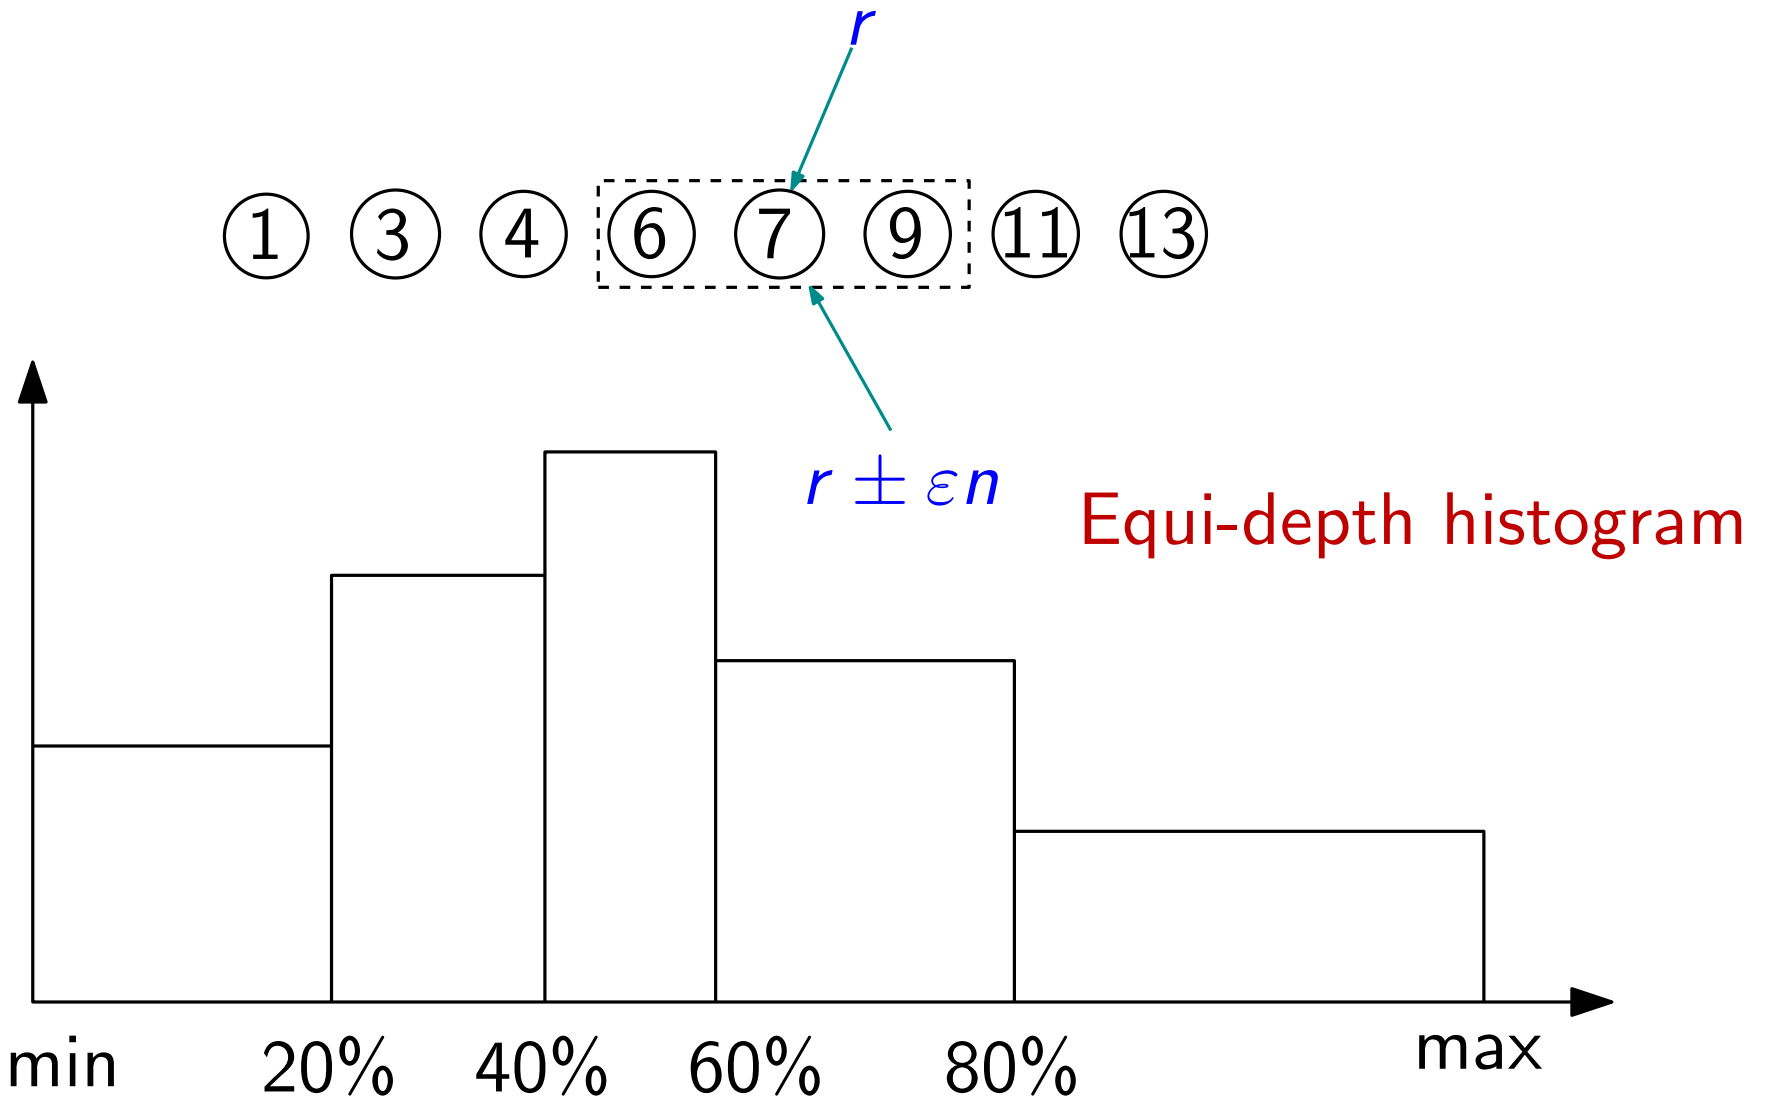
In a set of  $n$  values, the  $(r/n)$ -quantile is the value ranked at  $r$ .  
The 0.5-quantile is the **median**.



An  $\epsilon$ -approximate  $(r/n)$ -quantile is any value ranked between  $[r - \epsilon n, r + \epsilon n]$ .

# Quantiles

In a set of  $n$  values, the  $(r/n)$ -quantile is the value ranked at  $r$ .  
The 0.5-quantile is the **median**.



# Quantiles: Previous Solutions

- Sketching: Each node computes a sketch of its own data and sends it to the coordinator.

Sketch size:  $O(1/\varepsilon)$

Communication cost:  $O(k/\varepsilon)$

- Random sampling


Uniformly randomly sample a subset of size  $O(1/\varepsilon^2)$

- We can achieve:  $O(\sqrt{k}/\varepsilon)$

Typical values of  $\varepsilon = 10^{-3} \sim 10^{-6}$ ,  $k = 10^2 \sim 10^4$


We assume  $k < 1/\varepsilon^2$

# The Algorithm



The algorithm for each node

Sample each value with probability  $p$



Compute local ranks







# The Algorithm


At the base station:

Answering **value-to-rank** query

Given any value  $x$ , estimates its rank  $r(x)$



# The Algorithm




At the base station:


Answering **value-to-rank** query

Given any value  $x$ , estimates its rank  $r(x)$

predecessor  $r(10)?$



# The Algorithm




At the base station:

Answering **value-to-rank** query


Given any value  $x$ , estimates its rank  $r(x)$

predecessor

$r(10)?$



# The Algorithm




At the base station:

Answering **value-to-rank** query

Given any value  $x$ , estimates its rank  $r(x)$

predecessor

$r(10)?$



$$\hat{r}(10) = 5 + 2/p$$

# Correctness

Will show:  $\hat{r}(x)$  is an unbiased estimator of  $r(x)$  with standard deviation  $\varepsilon n$ .

$r(10)?$

1

5

6

10

11

15

# Correctness

Will show:  $\hat{r}(x)$  is an unbiased estimator of  $r(x)$  with standard deviation  $\varepsilon n$ .

$r(10)$ ?

5

15

# Correctness

Will show:  $\hat{r}(x)$  is an unbiased estimator of  $r(x)$  with standard deviation  $\varepsilon n$ .

$r(10)?$

5

?

15

# Correctness

Will show:  $\hat{r}(x)$  is an unbiased estimator of  $r(x)$  with standard deviation  $\varepsilon n$ .

$r(10)?$

5

?

15



Follows a geometric distribution (almost)

$$E[?] = 1/p$$

$$\text{Var}[?] \leq 1/p^2$$



# Correctness

Will show:  $\hat{r}(x)$  is an unbiased estimator of  $r(x)$  with standard deviation  $\varepsilon n$ .

$r(10)?$

5

?

15

Follows a geometric distribution (almost)

$$E[?] = 1/p$$

$$\text{Var}[?] \leq 1/p^2$$

$$\text{Set } p = \frac{\sqrt{k}}{\varepsilon n}$$


$$\text{Var}[\hat{r}(x)] \leq k/p^2 = (\varepsilon n)^2$$

Total cost:  $np = \sqrt{k}/\varepsilon$  in expectation

# Outline


- Model of computation
- Frequency estimation (heavy hitters)
- Quantiles (order statistics)
- **Other problems**

# $\varepsilon$ -approximate range counting



Let  $P$  be a set of  $n$  points in the plane. Compute a summary structure so that, for any range  $Q$  (from a certain range space),  $|P \cap Q|$  can be extracted with error  $\varepsilon n$


# $\epsilon$ -approximations



$S \subseteq P$  is an  $\epsilon$ -approximation of  $P$  if for any  $Q$  (from a certain range space),


$$|P \cap Q| = |S \cap Q| \cdot \frac{n}{|S|} \pm \epsilon n$$

# $\varepsilon$ -approximations




Size of  $\varepsilon$ -approximations

# $\epsilon$ -approximations over $k$ distributed data sets



$$\sqrt{k} \cdot 1/\epsilon \log^{O(1)}(1/\epsilon)$$



$$k^{1/3} \cdot 1/\epsilon^{4/3}$$

# The General Question

For what problems can we do better than  $k \times$  sketch size?

# The General Question

For what problems can we do better than  $k \times$  sketch size?

- Some positive results in this talk
- Negative results in Qin Zhang's talk
  - Number of distinct elements
  - Frequency moments



# References

- Optimal Sampling Algorithms for Frequency Estimation in Distributed Data. Zengfeng Huang, Ke Yi, Yunhao Liu, Guihai Chen. INFOCOM 2011.
- Sampling Based Algorithms for Quantile Computation in Sensor Networks. Zengfeng Huang, Lu Wang, Ke Yi, and Yunhao Liu. SIGMOD 2011.