Randomized Algorithms for Tracking Distributed Count, Frequencies, and Ranks

Zengfeng Huang, Ke Yi
HKUST

Qin Zhang
MADALGO

PODS'12

The Model

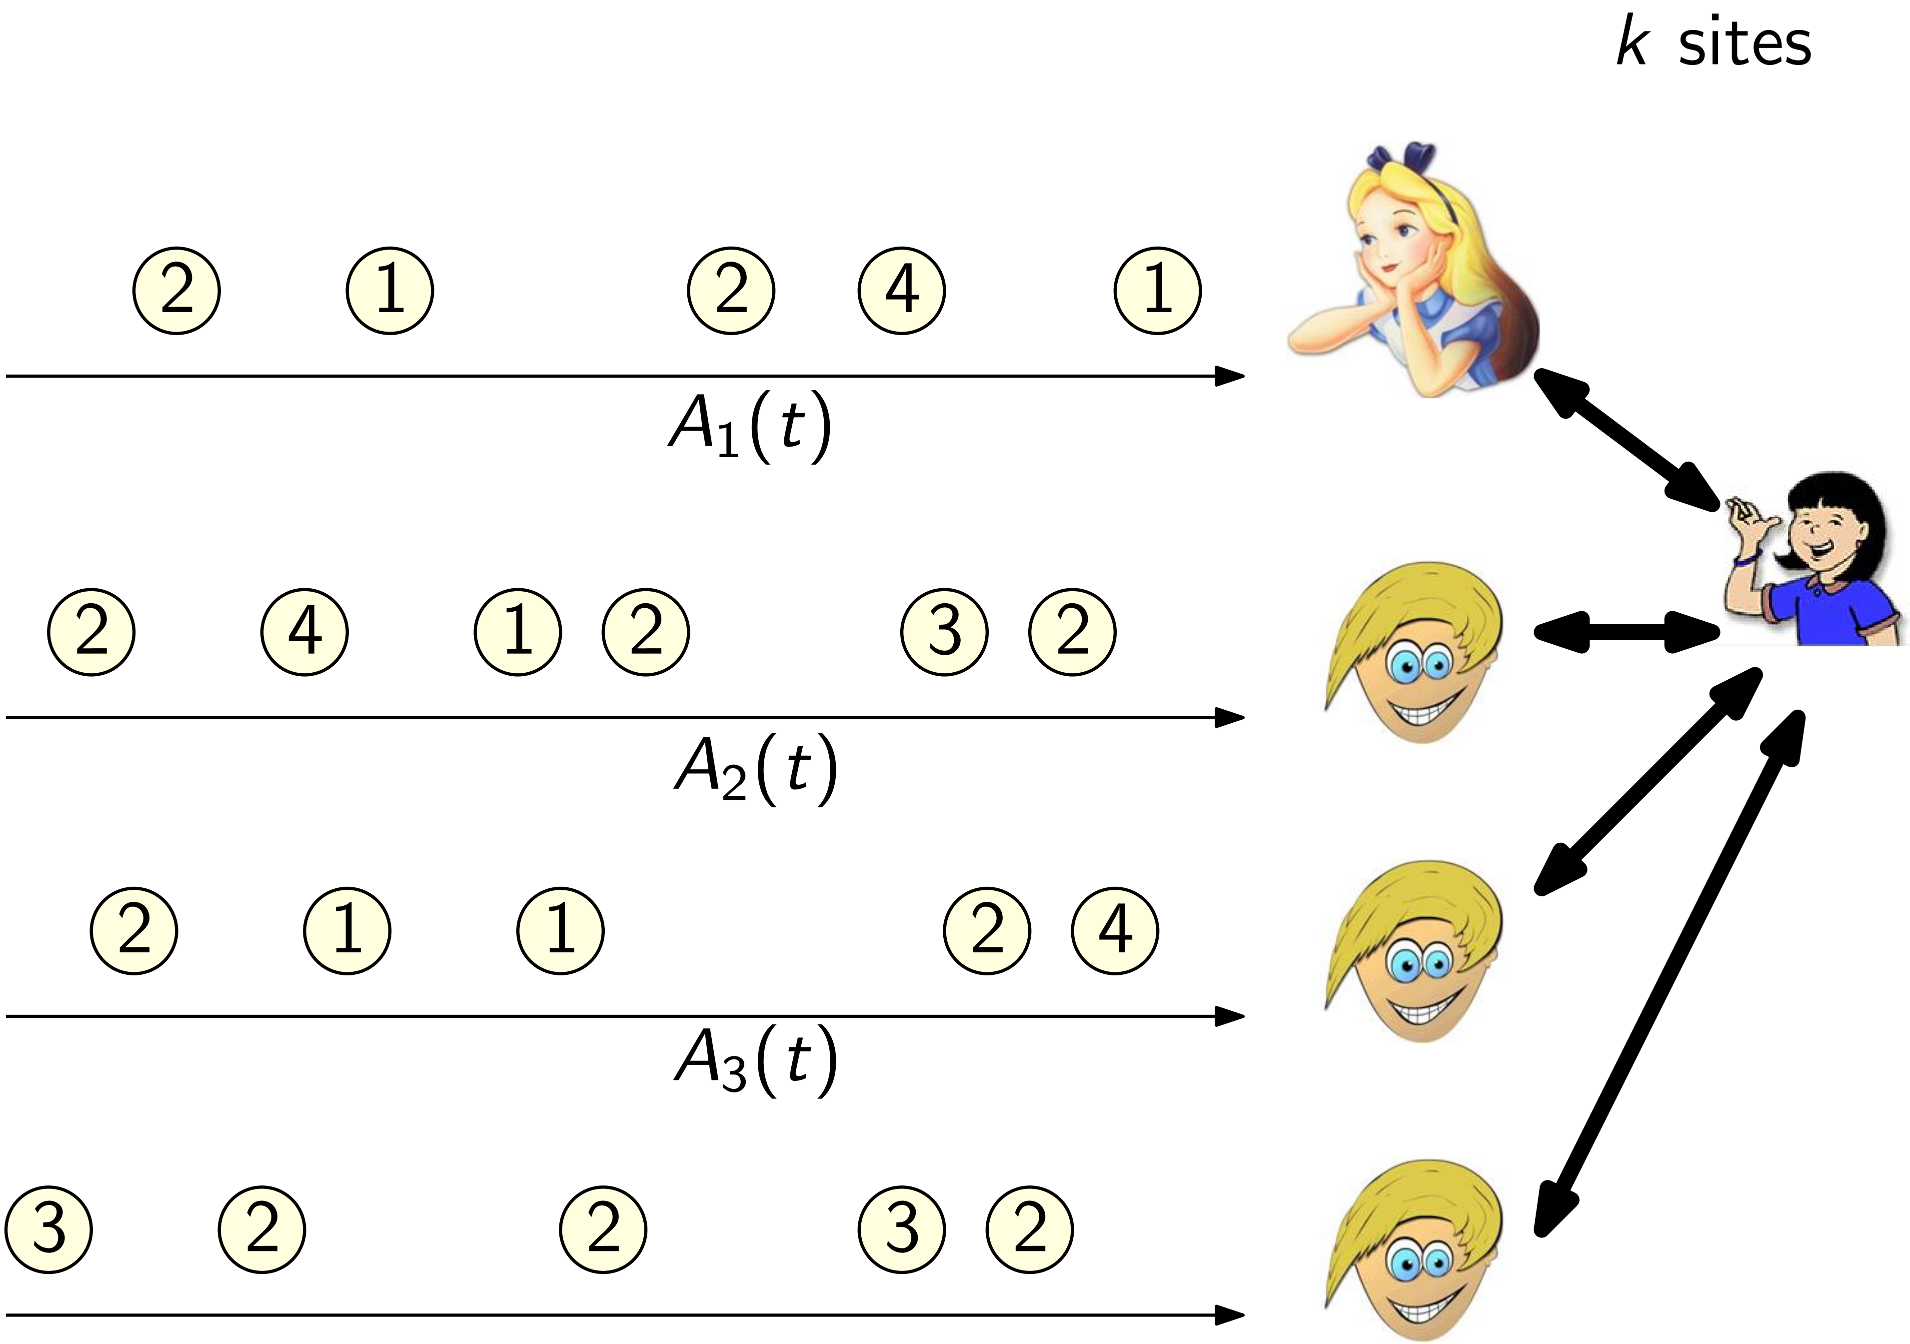

The Model

Coordinator tries to compute $f(A_1(t) \uplus A_2(t) \cdots A_k(t))$ for all t

k sites

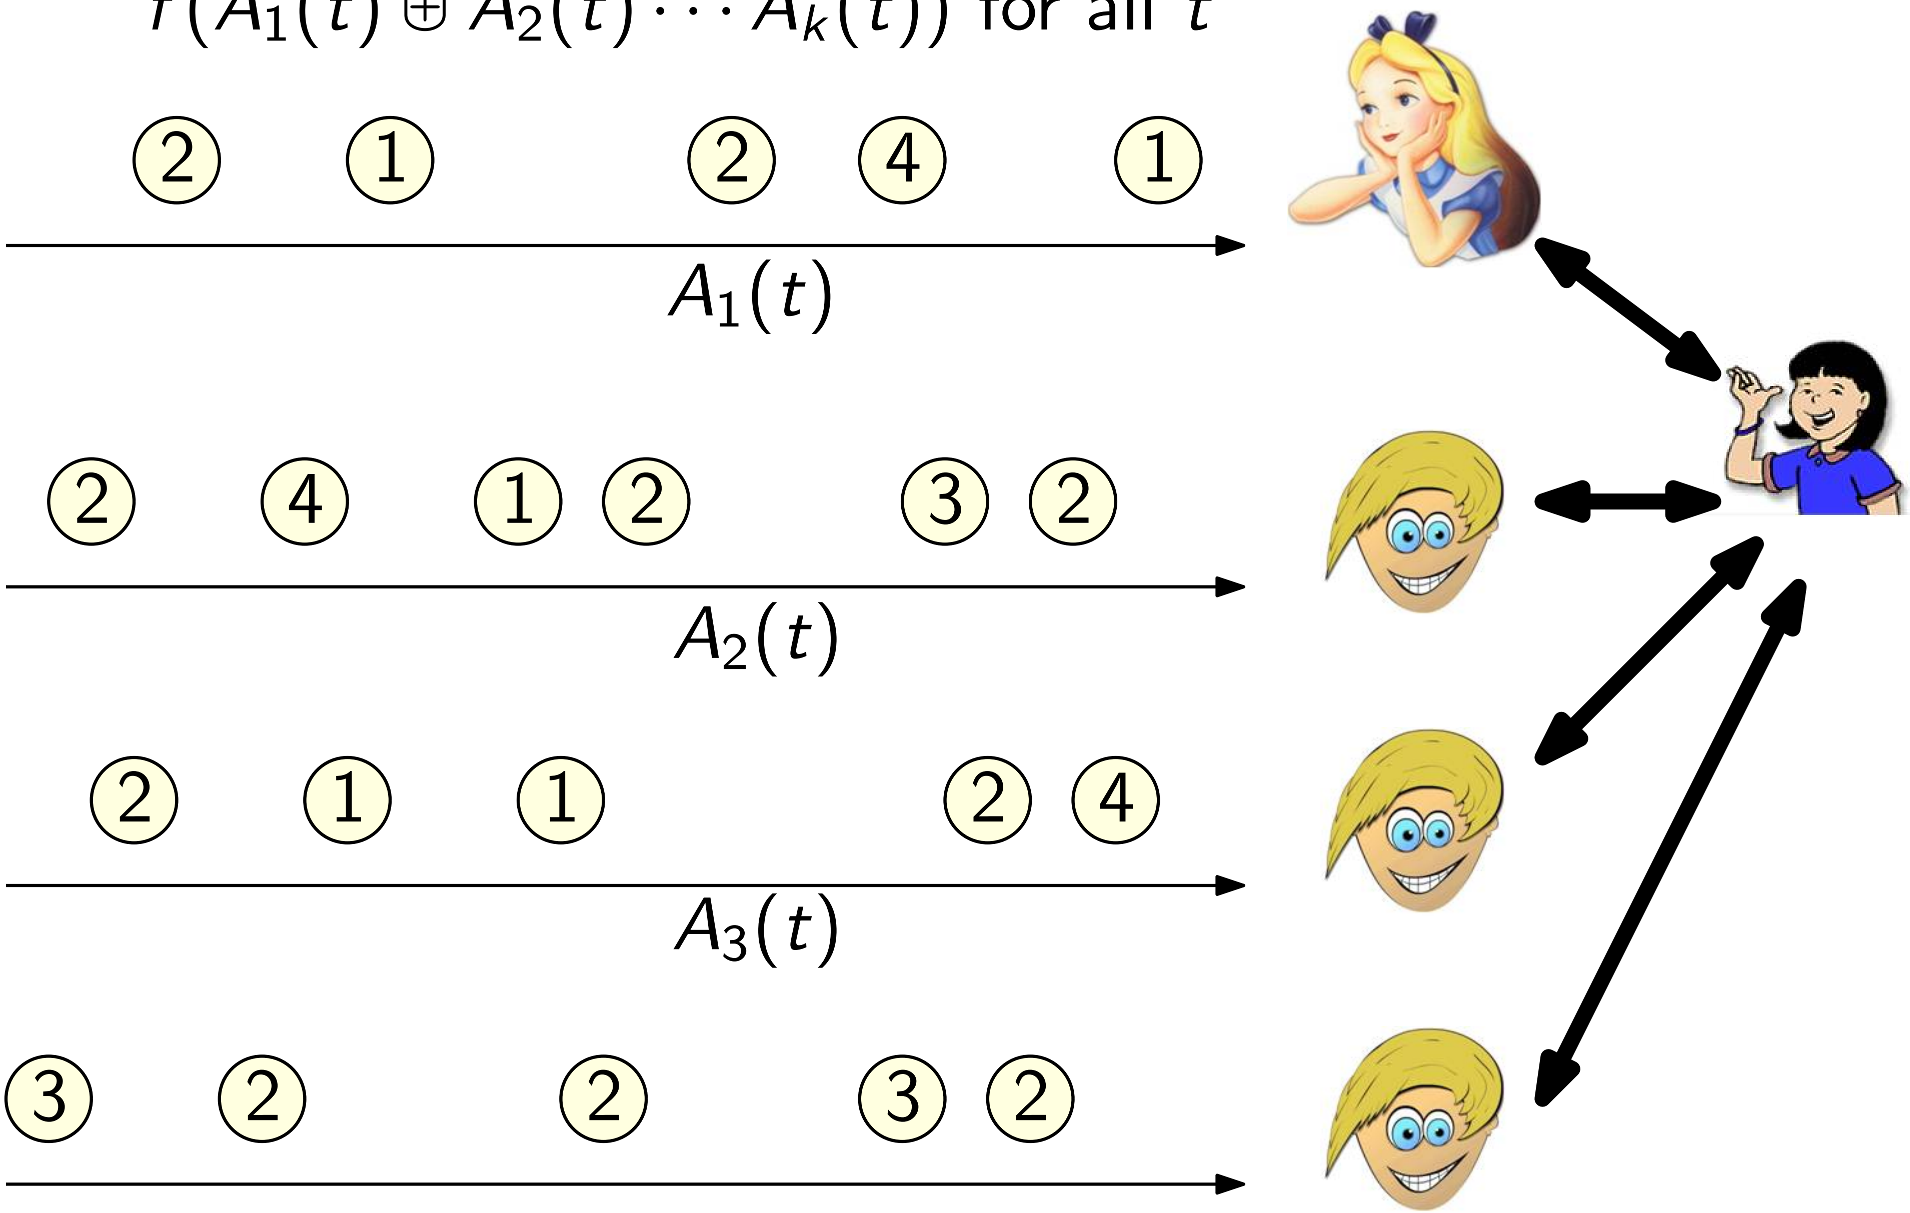

Generalization of Two Models

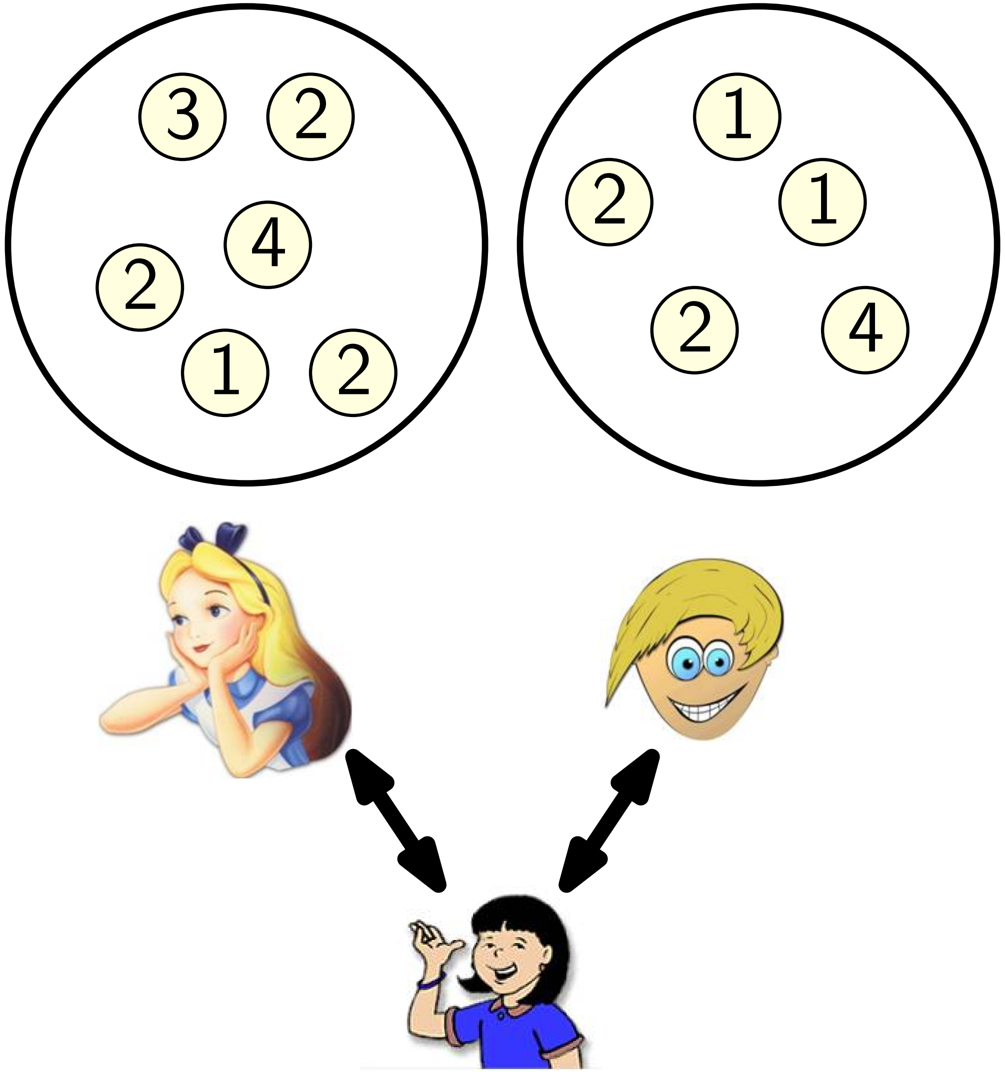

Communication model
(One-shot model)

Generalization of Two Models

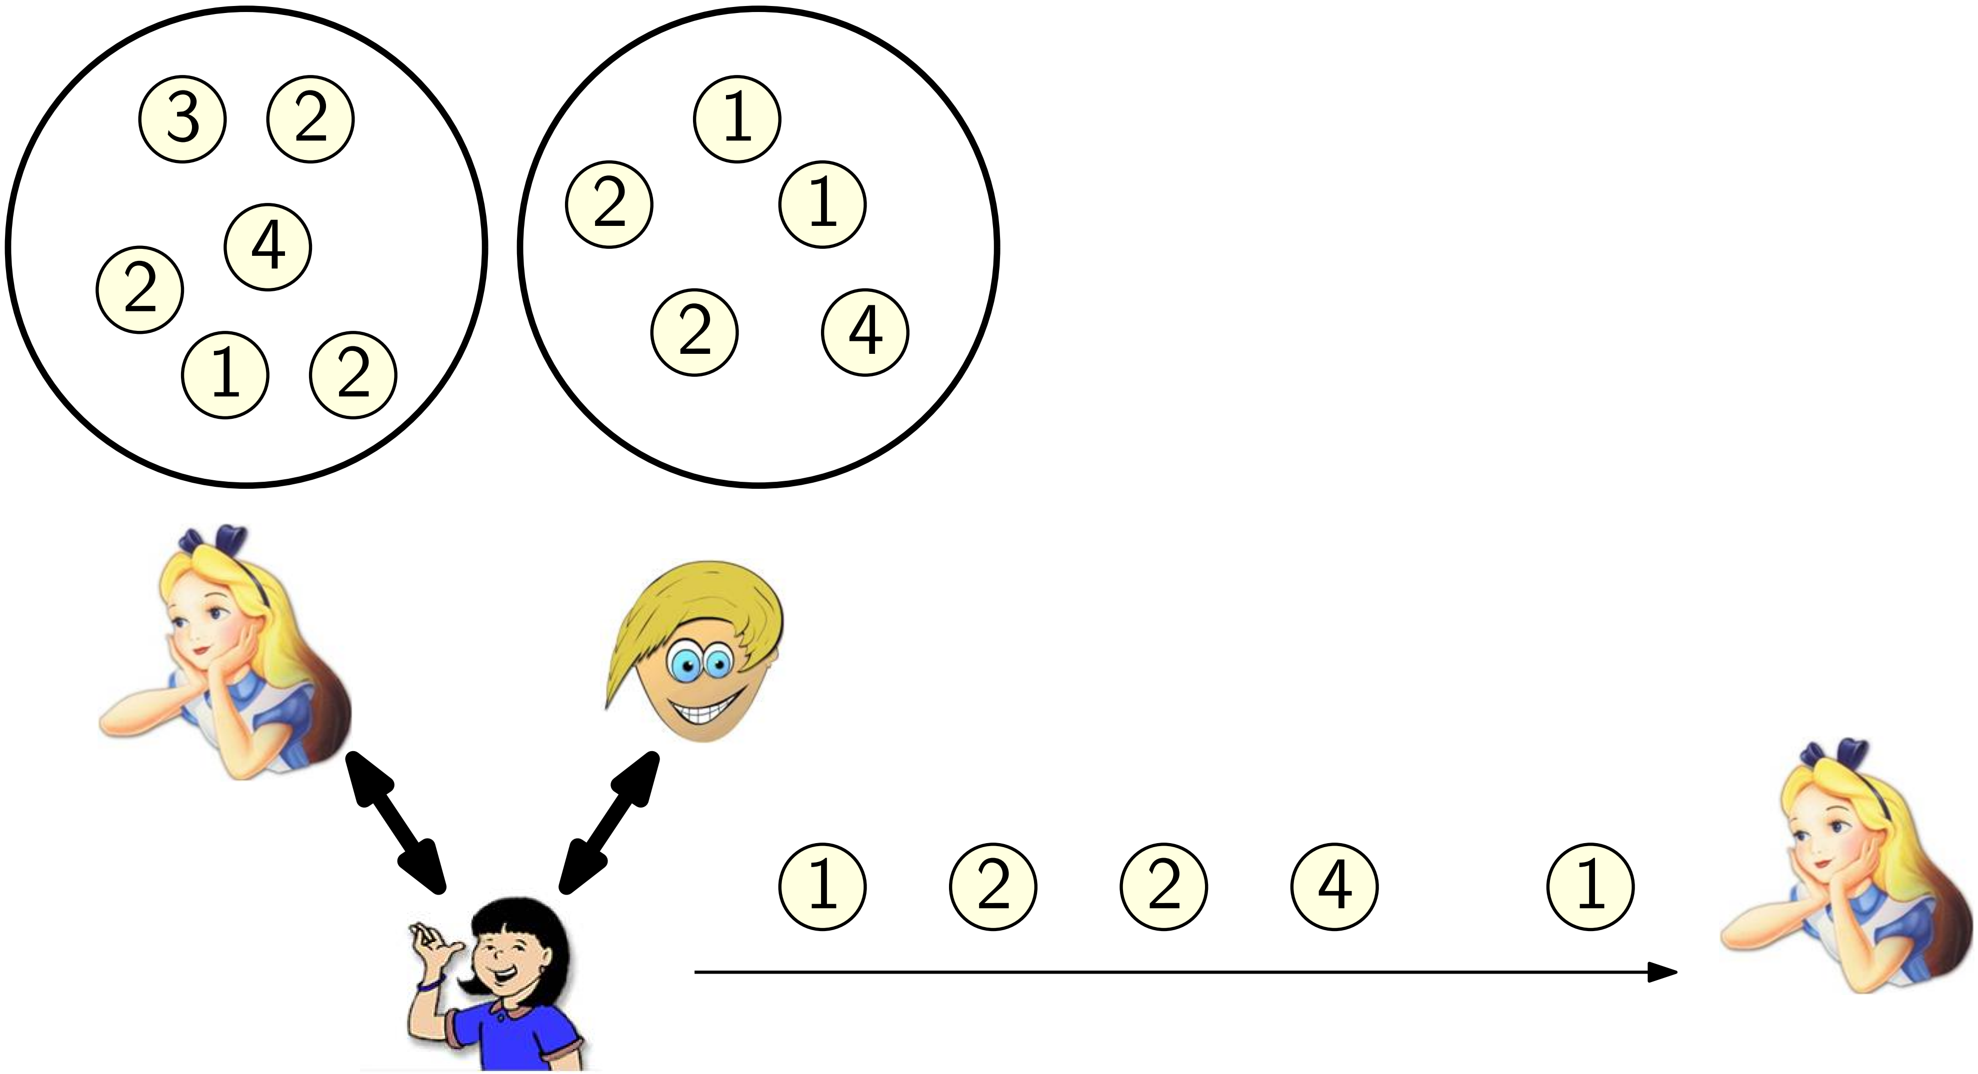

Communication model
(One-shot model)

Data stream model

Generalization of Two Models

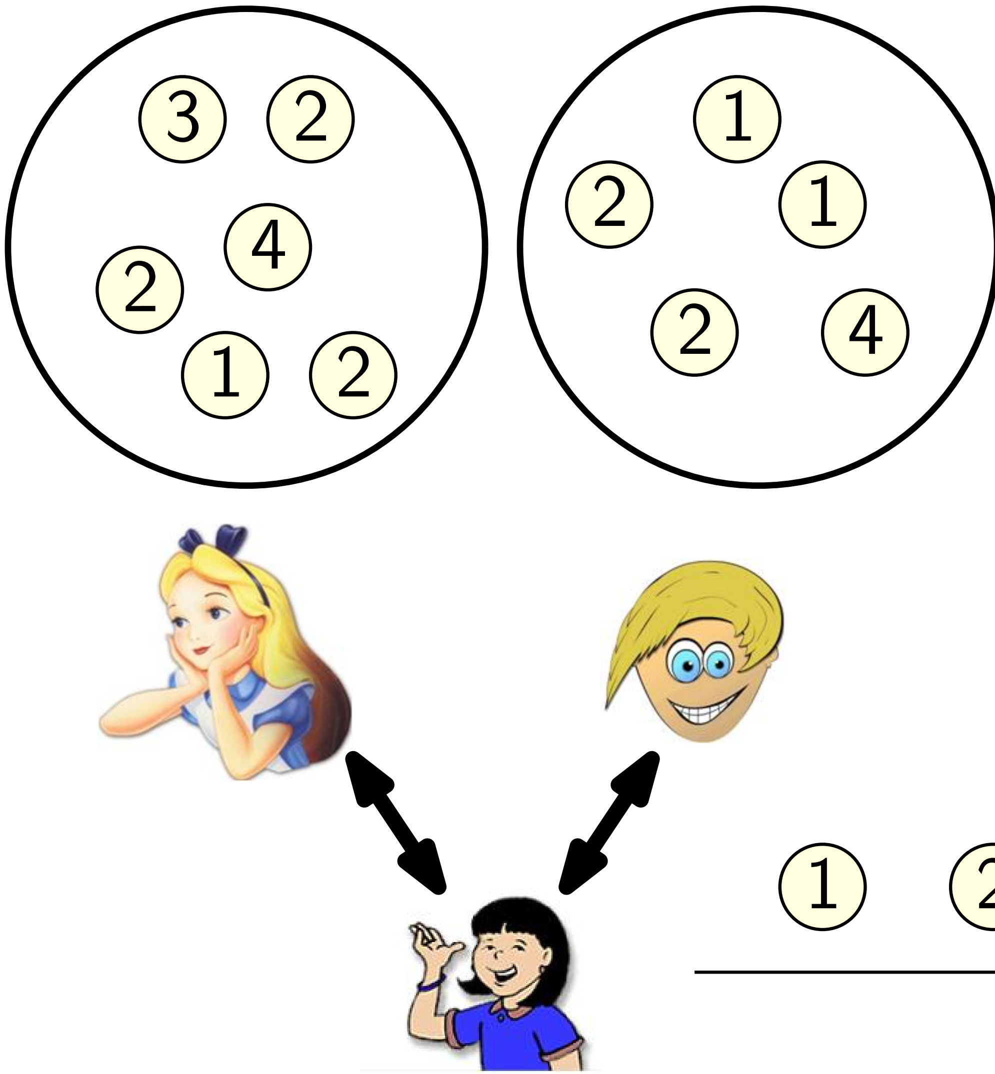

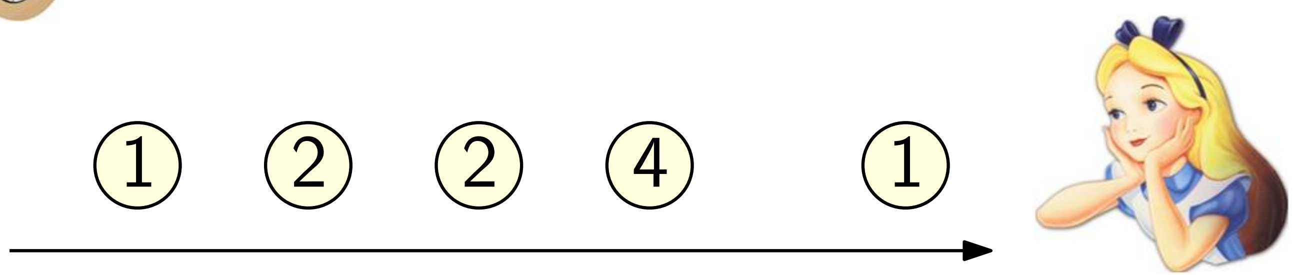

Goal

- Communication cost
- Space

Communication model
(One-shot model)

Data stream model

Warm Up: Count tracking

$f(A) = |A(t)|$ (Trivial in previous models)

Warm Up: Count tracking

$f(A) = |A(t)|$ (Trivial in previous models)

Let $n_i = A_i(t)$ be the local count

Track $n = \sum_i n_i$ continously within additive error εn at any time

Count tracking

Communication: $O(k/\varepsilon \cdot \log n)$
Space per site: $O(1)$
Tight

[Yi, Zhang, PODS'09]

Count tracking

Communication: $O(k/\varepsilon \cdot \log n)$
Space per site: $O(1)$
Tight

[Yi, Zhang, PODS'09]

Communication: $O(\sqrt{k}/\varepsilon \cdot \log n)$
Space per site: $O(1)$
Tight

New

Deterministic

Each site sets some thresholds
Sends a message when n_i exceed a threshold

[Yi, Zhang, PODS'09]

Deterministic

Each site sets some thresholds

Sends a message when n_i exceed a threshold

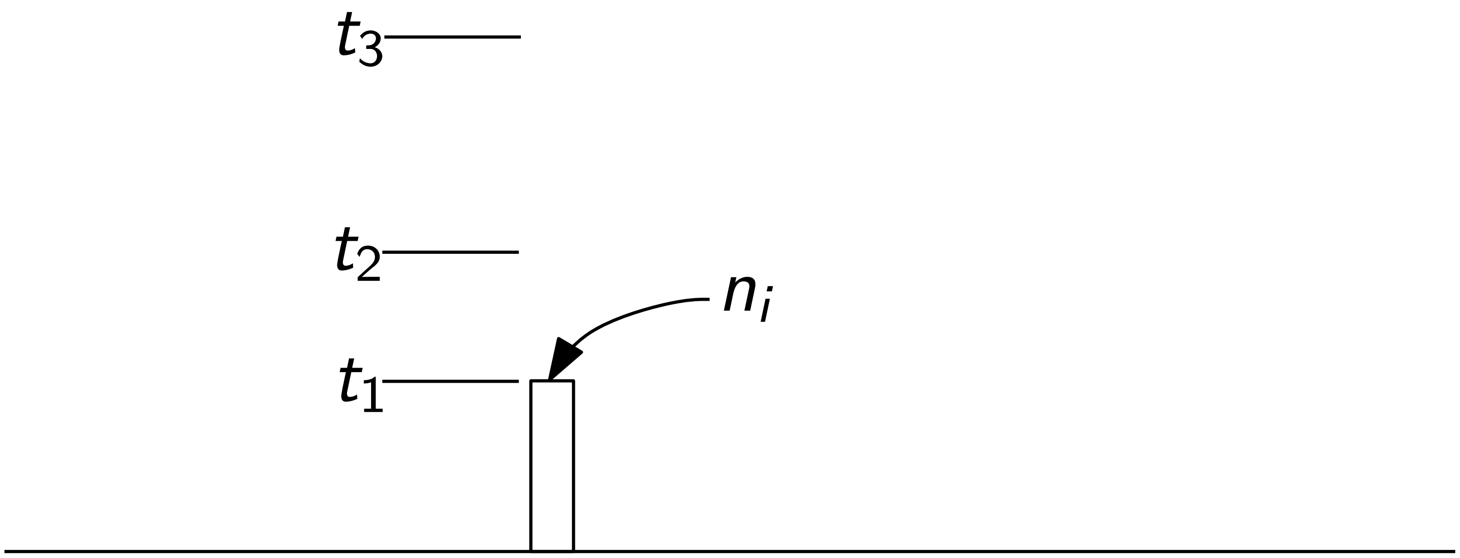

[Yi, Zhang, PODS'09]

Deterministic

Each site sets some thresholds

Sends a message when n_i exceed a threshold

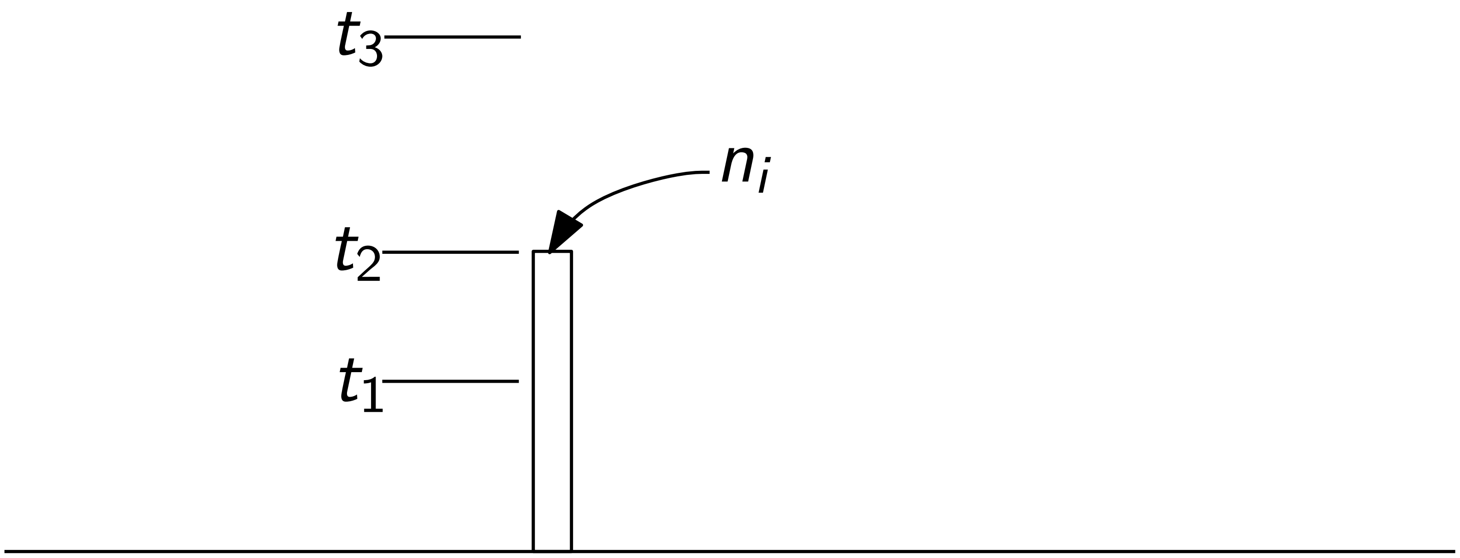

[Yi, Zhang, PODS'09]

Deterministic

Each site sets some thresholds
Sends a message when n_i exceed a threshold

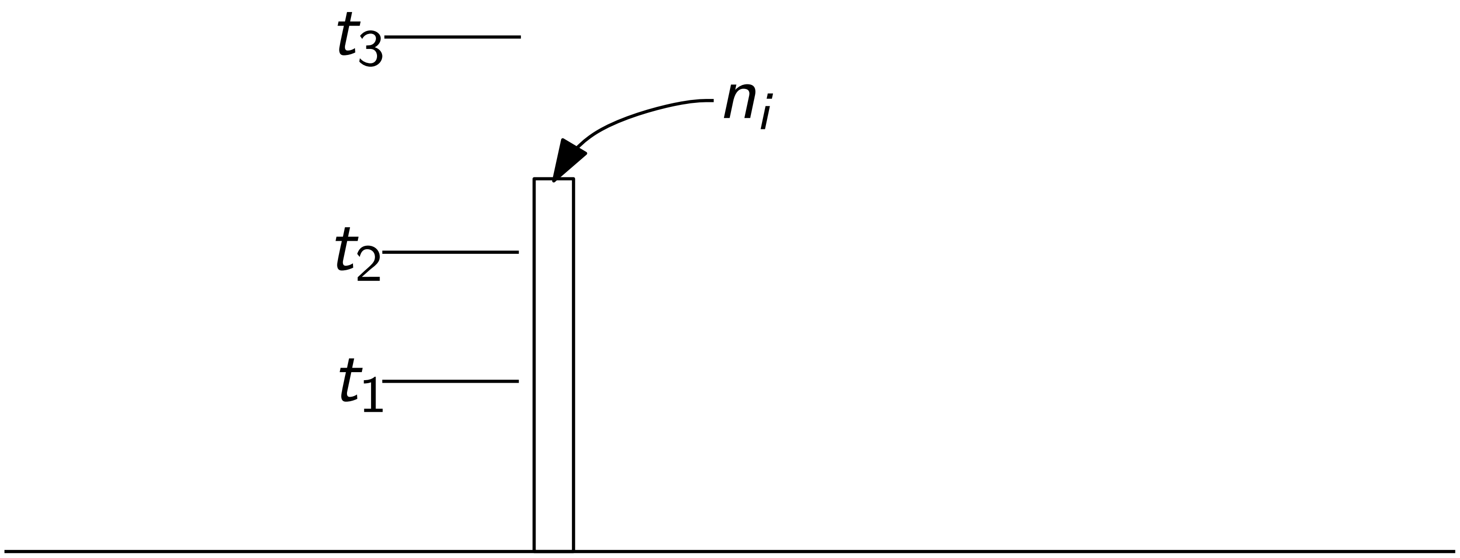

[Yi, Zhang, PODS'09]

Deterministic

Each site sets some thresholds
Sends a message when n_i exceed a threshold

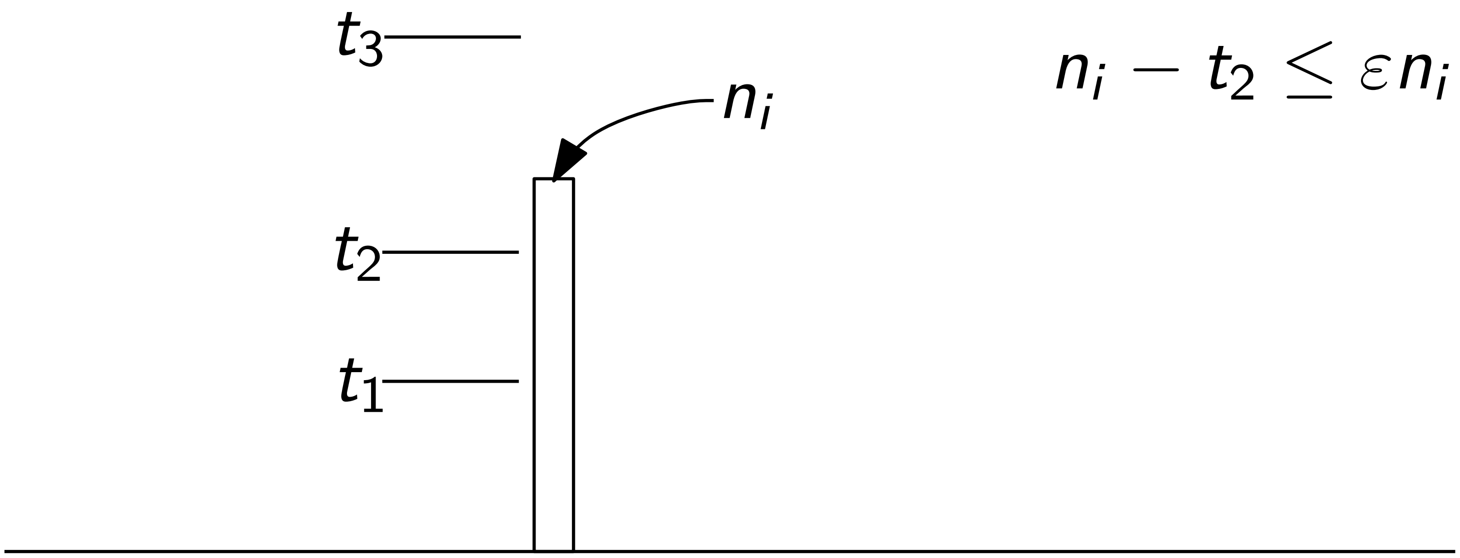

$$n_i - t_2 \leq \varepsilon n_i$$

[Yi, Zhang, PODS'09]

Randomized algorithm

Make decision based on random bits

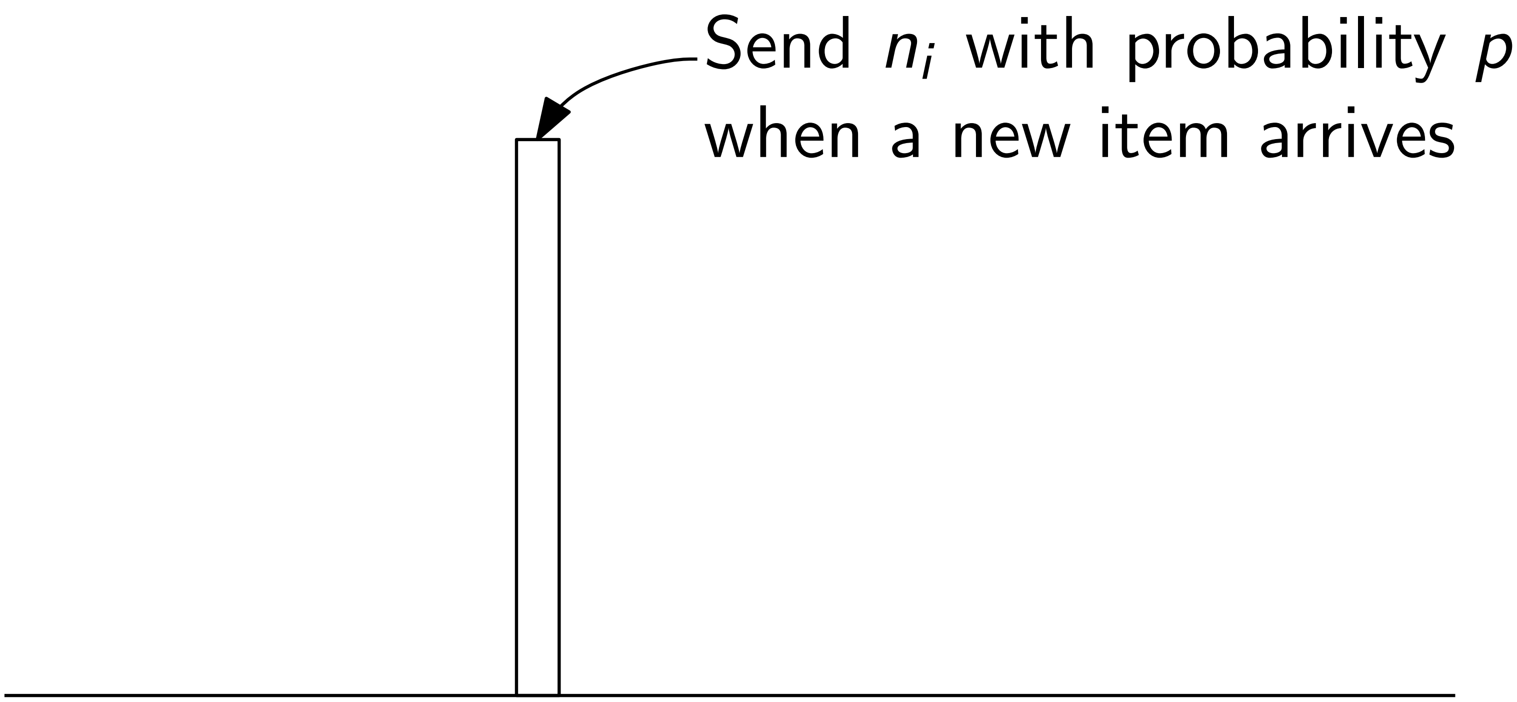

Analysis

$n_i - \bar{n}_i$ is a random variable

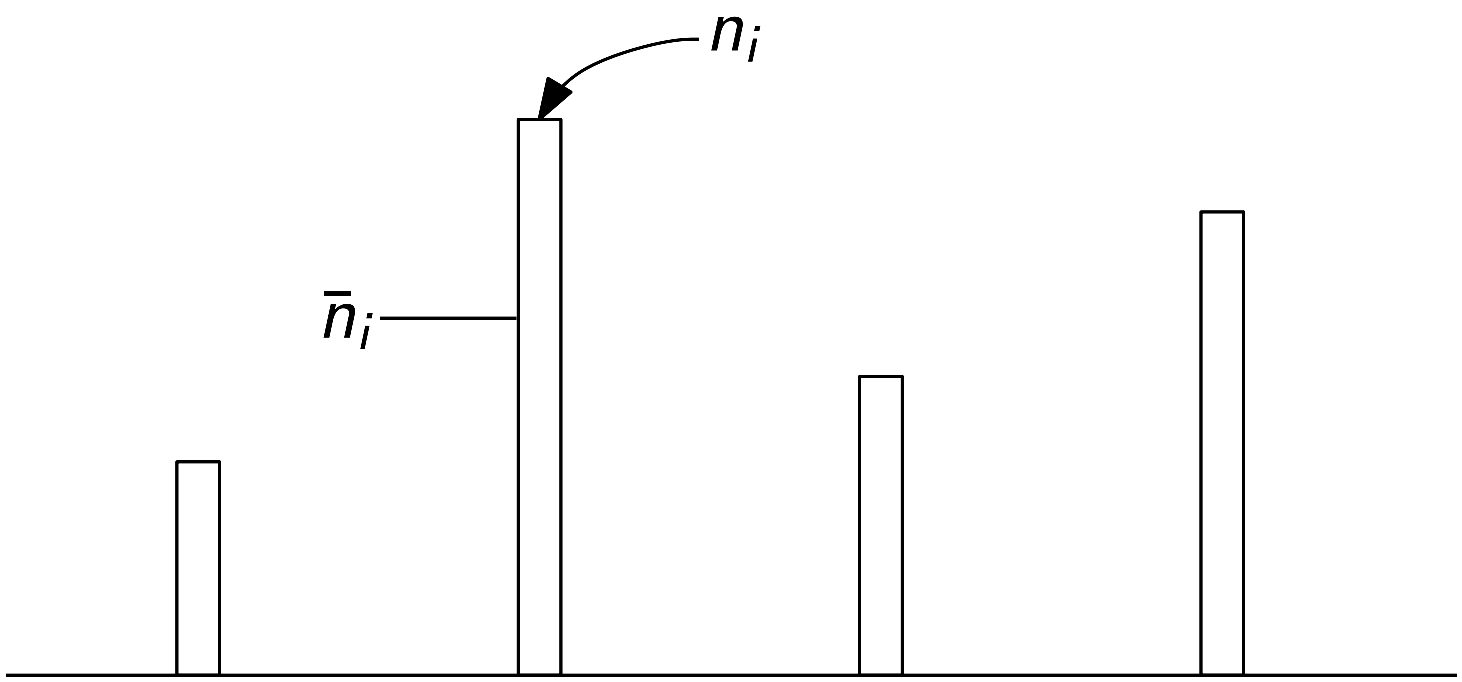

Analysis

$n_i - \bar{n}_i$ is a random variable

$$\hat{n}_i = \begin{cases} \bar{n}_i - 1 + 1/p, & \text{if } t > 0; \\ 0, & \text{else.} \end{cases}$$

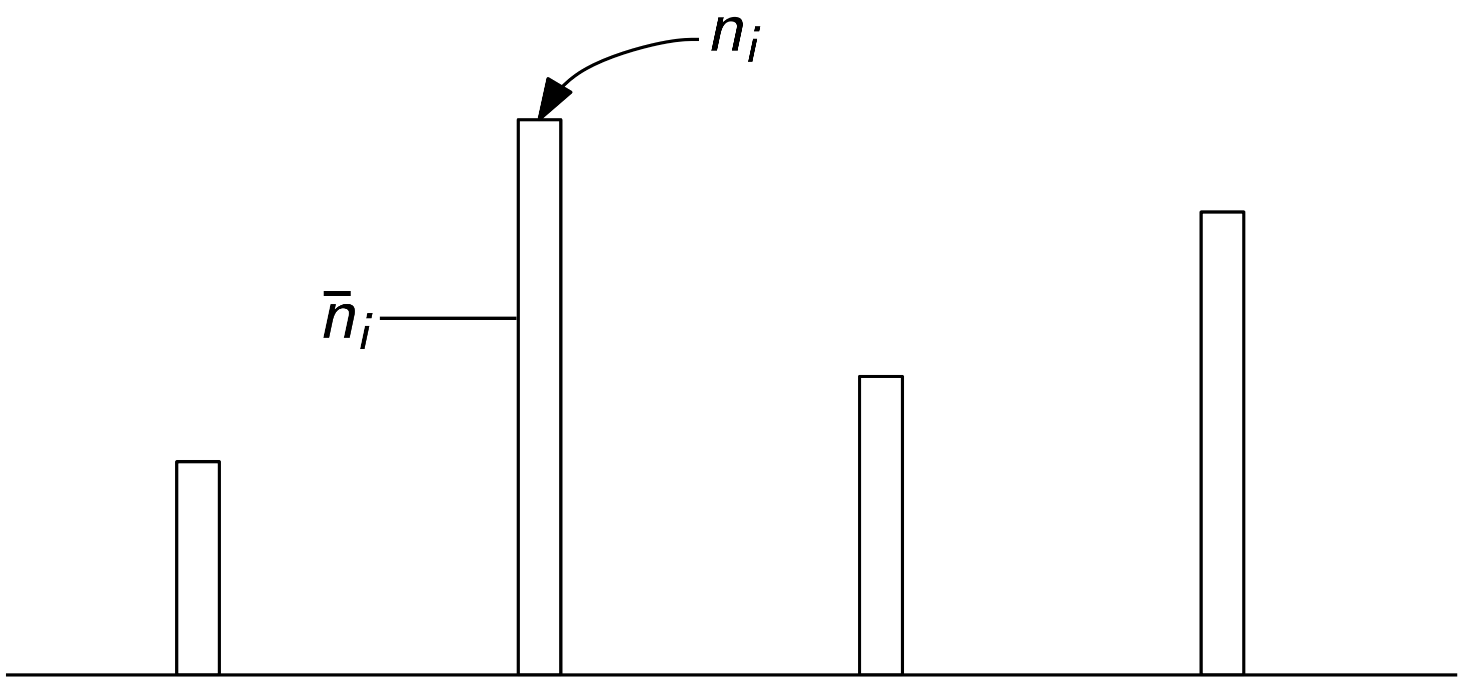

Analysis

$n_i - \bar{n}_i$ is a random variable

$$\hat{n}_i = \begin{cases} \bar{n}_i - 1 + 1/p, & \text{if } t > 0; \\ 0, & \text{else.} \end{cases}$$

$E[\hat{n}_i] = n_i$, $\mathrm{Var}[\hat{n}_i] = 1/p^2$

Analysis

$n_i - \bar{n}_i$ is a random variable

$$\hat{n}_i = \begin{cases} \bar{n}_i - 1 + 1/p, & \text{if } t > 0; \\ 0, & \text{else.} \end{cases}$$

$\mathrm{E}[\hat{n}_i] = n_i$, $\mathrm{Var}[\hat{n}_i] = 1/p^2$

$$\hat{n} = \sum \hat{n}_i$$

$\mathrm{E}[\hat{n}] = \sum \hat{n}_i = n$, $\mathrm{Var}[\hat{n}] = k/p^2$

Rounds

Chebyshev

SD less than $\varepsilon n \rightarrow p = O(\sqrt{k}/\varepsilon n)$
constant probability of success

Rounds

Chebyshev

SD less than $\varepsilon n \to p = O(\sqrt{k}/\varepsilon n)$
constant probability of success

- Track a 2-approximation $\bar{n}$ of n
  - Broadcast $\bar{n}$ whenever $\bar{n}$ doubles
  - Set $p = \frac{\sqrt{k}}{2\bar{n}}$
- Divide the tracking period into rounds
  - n approximately doubles in a round
  - p is fixed in a round

Rounds

- Communication cost
  - Tracking a constant approximation $O(k \log n)$
  - number of messages in a round: $O(np) = O(\sqrt{k}/\varepsilon)$
  - Total: $O(k \log n + \sqrt{k}/\varepsilon \cdot \log n)$

Randomized lower bound

One-way communication lower bound: $\Omega(k/\varepsilon \cdot \log n)$

Randomized lower bound

One-way communication lower bound: $\Omega(k/\varepsilon \cdot \log n)$

- No global information
- Can not distinguish 2 extremes:
  1. Evenly distributed
  2. All the items arrive at one site

Randomized lower bound

Communication lower bound: $\Omega(\sqrt{k}/\varepsilon \cdot \log n)$

Randomized lower bound

Communication lower bound: $\Omega(\sqrt{k}/\varepsilon \cdot \log n)$

1-BIT problem

Set s to $k/2 + \sqrt{k}$ or $k/2 - \sqrt{k}$ randomly
Randomly pick s sites with input 1, others 0

Goal: determine s

Randomized lower bound

Lemma

Any deterministic algorithm that solves 1-bit problem has communication cost $\Omega(k)$

Randomized lower bound

Lemma

Any deterministic algorithm that solves 1-bit problem has communication cost $\Omega(k)$

Main idea:

- Communicate less than $o(k)$ bits, the uncertainty is still high
- The algorithm has good chance to make a mistake

Randomized lower bound

Hard input in round i

Divide the round into $\frac{1}{\varepsilon\sqrt{k}}$ subrounds

In any subround r:

- Set s to $k/2 + \sqrt{k}$ or $k/2 - \sqrt{k}$ randomly
- Randomly pick s sites, send them 2^i items each.

Randomized lower bound

Hard input in round i

Divide the round into $\frac{1}{\varepsilon\sqrt{k}}$ subrounds

In any subround r:

- Set s to $k/2 + \sqrt{k}$ or $k/2 - \sqrt{k}$ randomly
- Randomly pick s sites, send them 2^i items each.

Cost in a subround is $O(k)$:

- $n \leq O(\sqrt{k}/\varepsilon \cdot 2^i)$
- Error allowed: $O(\sqrt{k}2^i)$

Cost in a round is $O(\sqrt{k}/\varepsilon)$:

Frequent Items: Definition

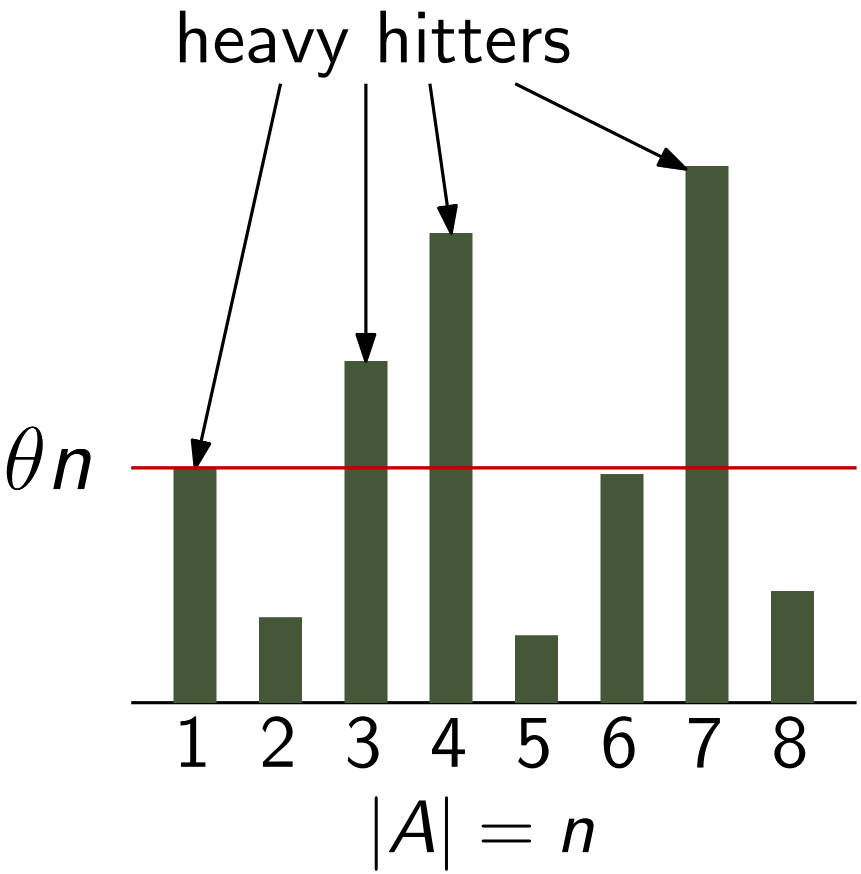

Frequent Items: Definition

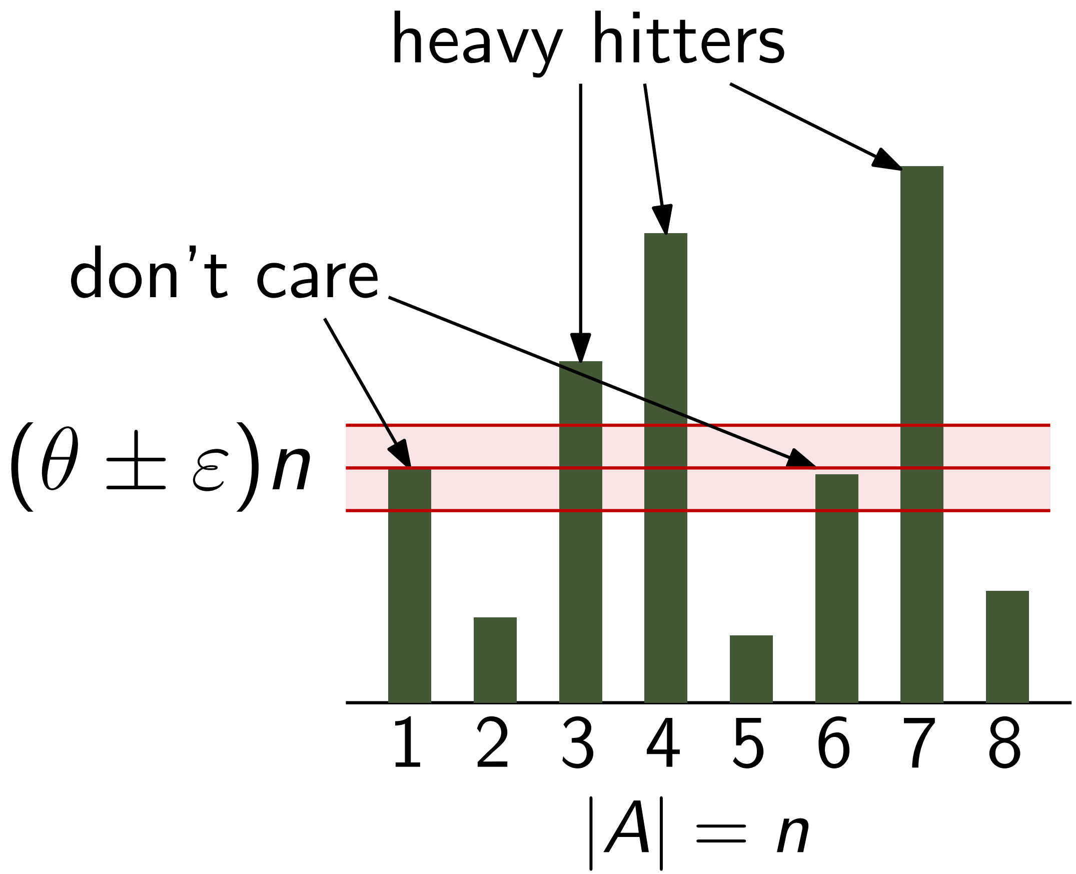

Frequent Items: Definition

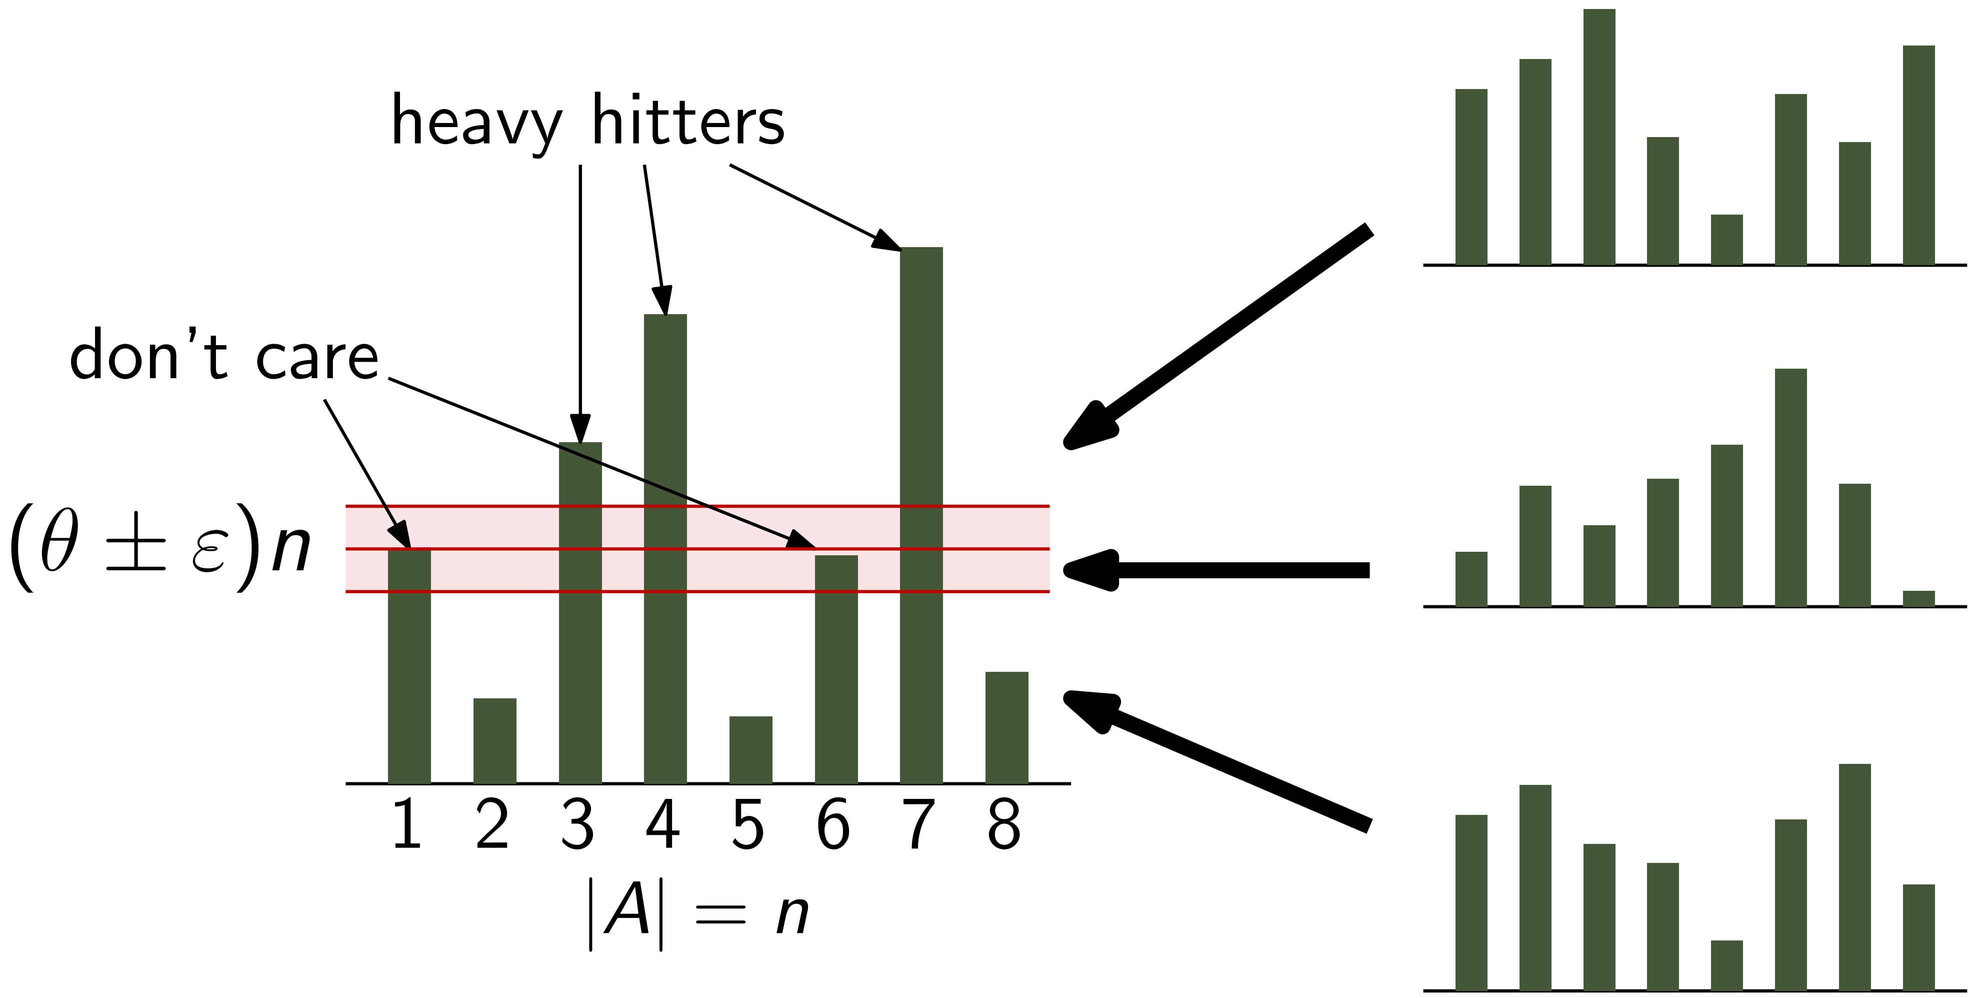

Frequent Items: Definition

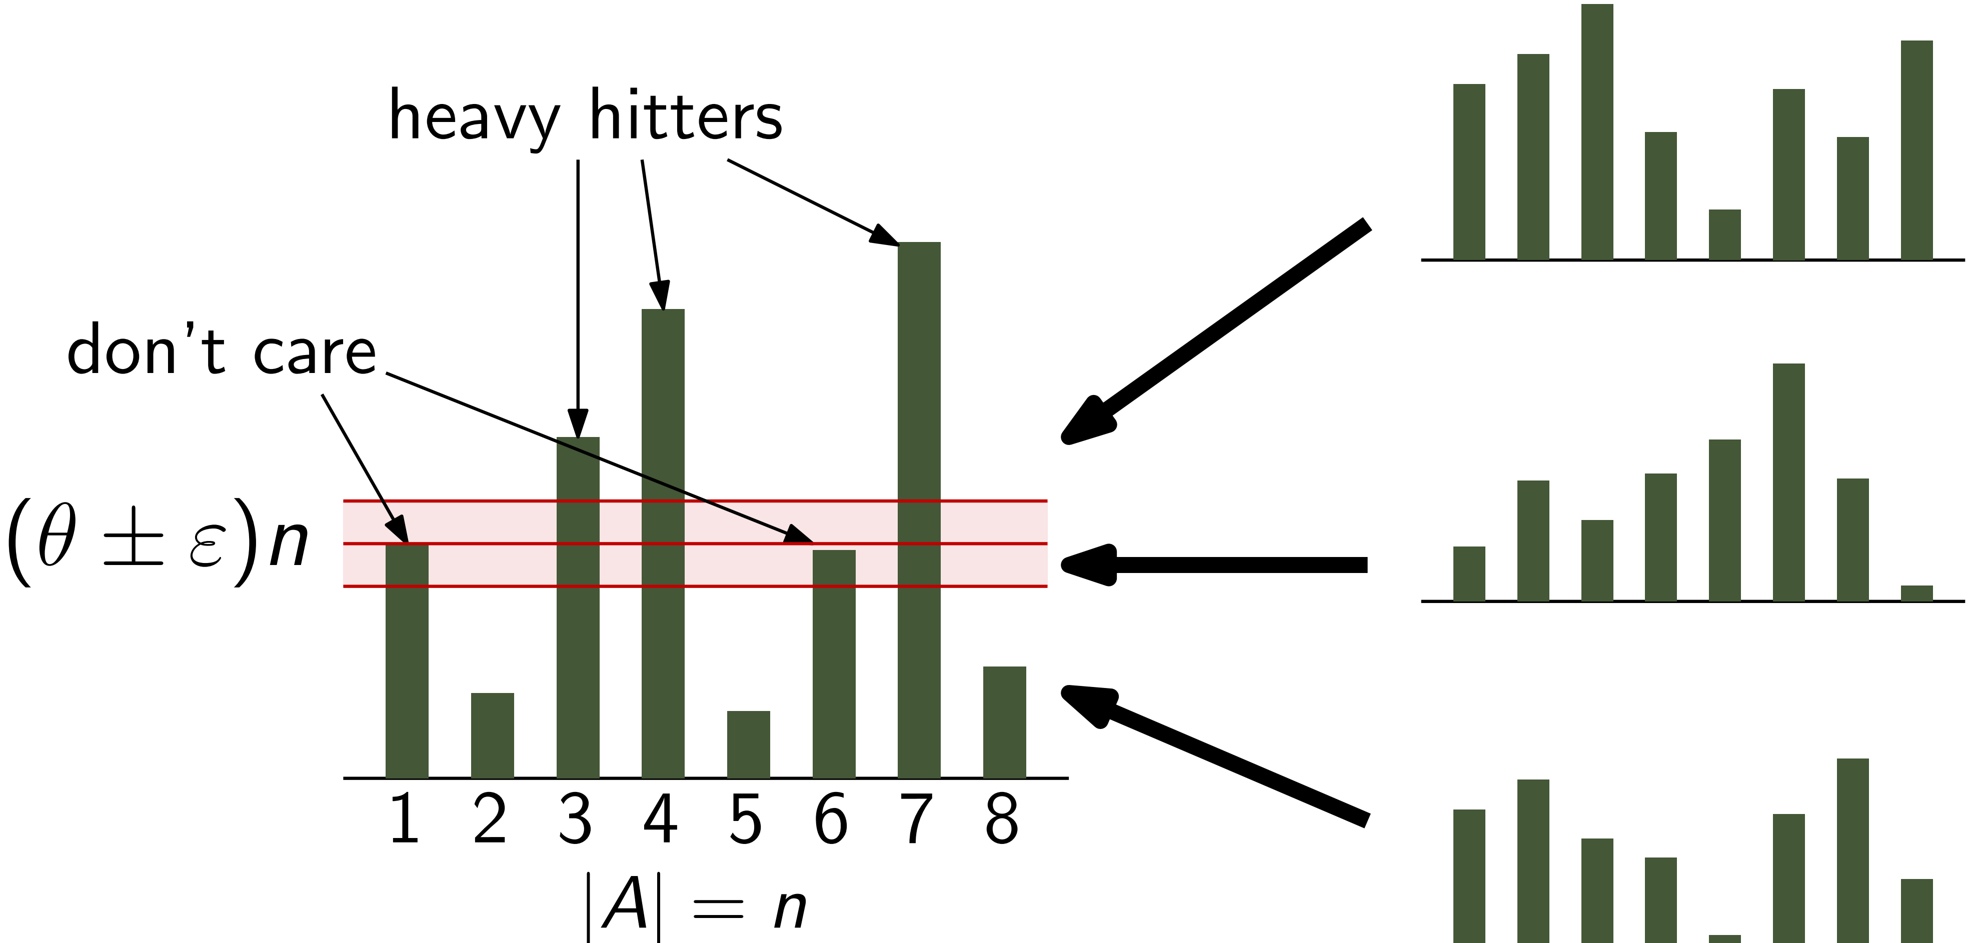

Frequency estimation with F_1 error

Estimate the frequency of every element with additive error εn.

Frequent Items: Algorithm in a round

Use the previous algorithm on each item i

- Maintain a count for each item at each site
- Space

Frequent Items: Algorithm in a round

Use the previous algorithm on each item i

- Maintain a count for each item at each site
- Space

Streaming algorithm

cost per site: $O(1/\varepsilon)$

- total: $O(k/\varepsilon)$
- improve to $O(\sqrt{k}/\varepsilon)$

Frequent Items: algorithm

Idea: maintain only large enough counts

Frequent Items: algorithm

Idea: maintain only large enough counts

i:

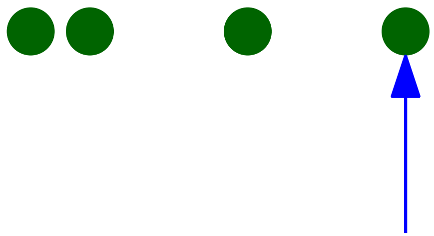

Start to count i with probability p

Frequent Items: algorithm

Idea: maintain only large enough counts

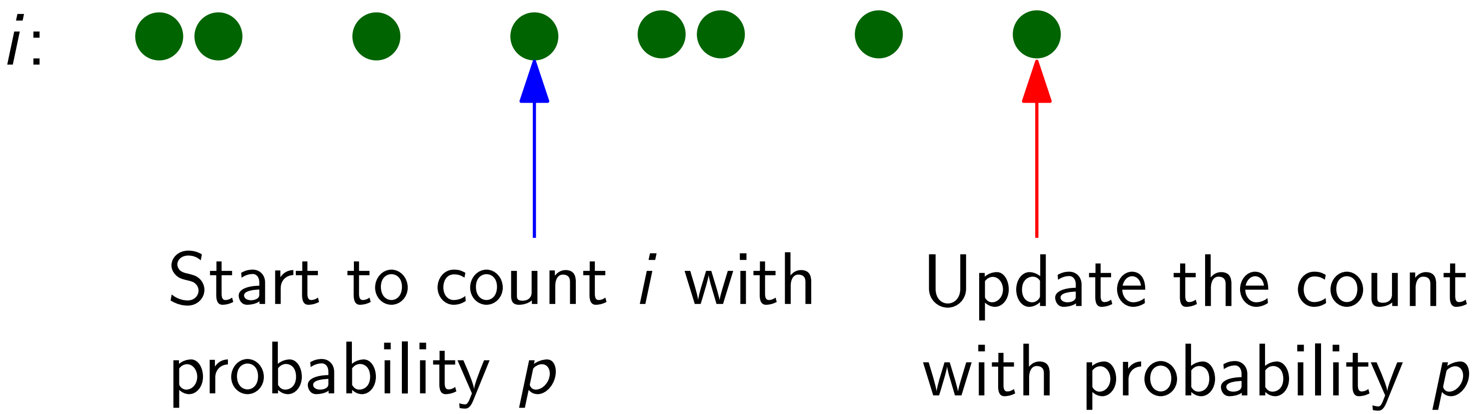

Frequent Items: analysis

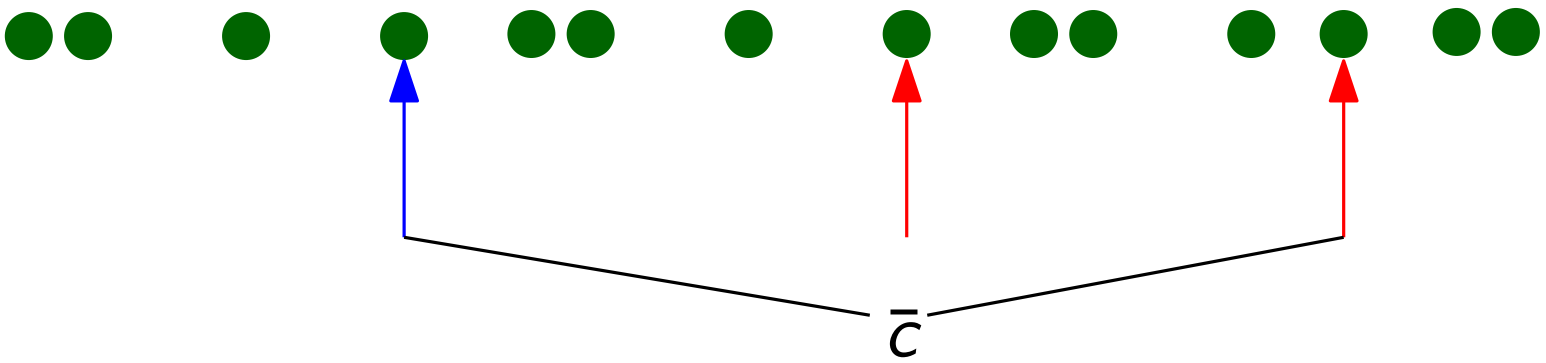

Coordinator only know $\bar{c}$

Frequent Items: analysis

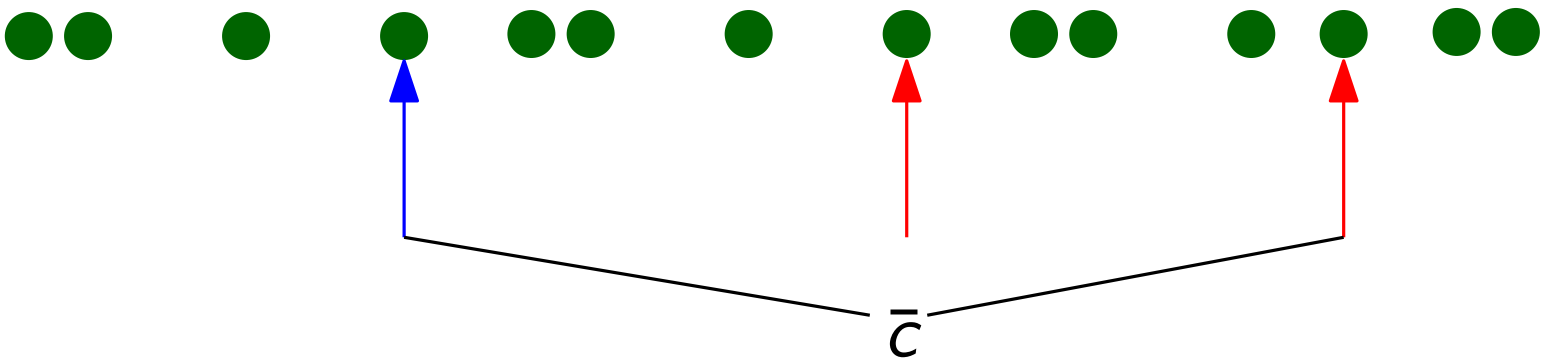

Coordinator only know $\bar{c}$

$$\hat{f}_i = \begin{cases} \bar{c} - 1 + 2/p, & \text{if } \bar{c} > 0; \\ 0, & \text{else.} \end{cases}$$

Frequent Items: analysis

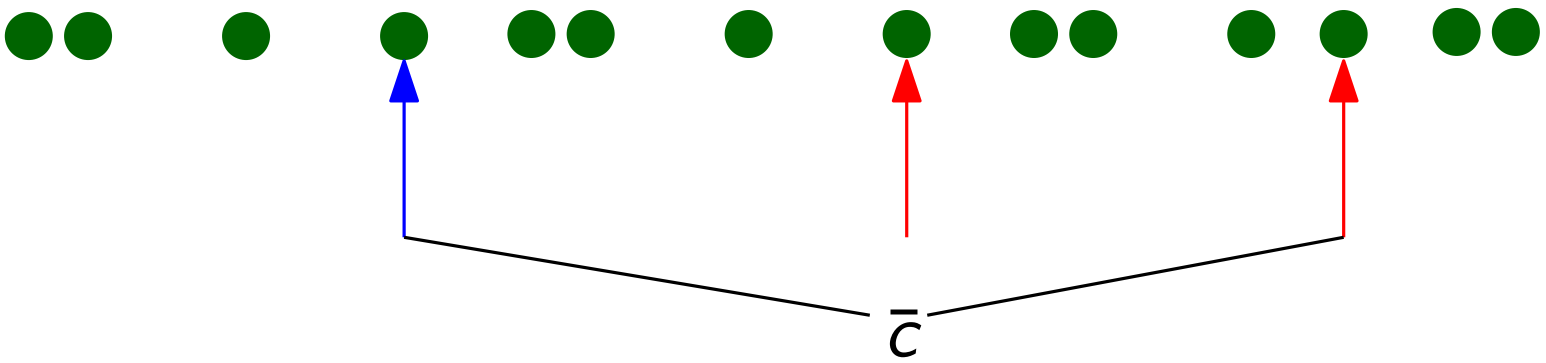

Coordinator only know $\bar{c}$

$$\hat{f}_i = \begin{cases} \bar{c} - 1 + 2/p, & \text{if } \bar{c} > 0; \\ 0, & \text{else.} \end{cases}$$

Bias might be as large as $\varepsilon n / \sqrt{k}$

Frequent Items: analysis

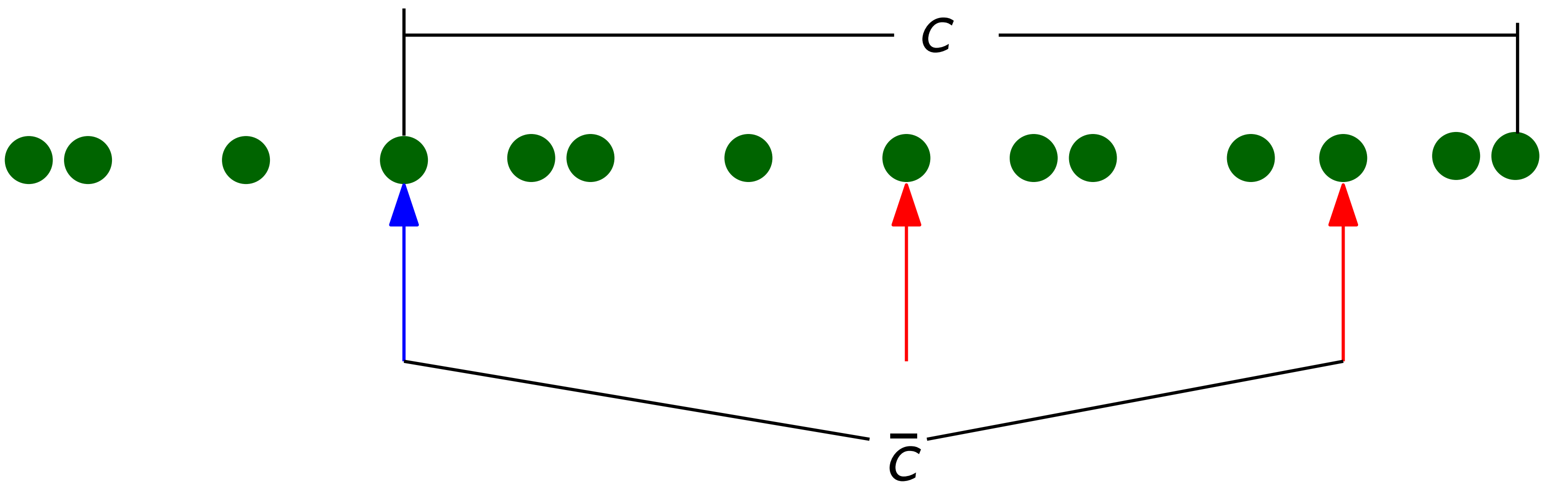

Coordinator only know $\bar{c}$

Frequent Items: analysis

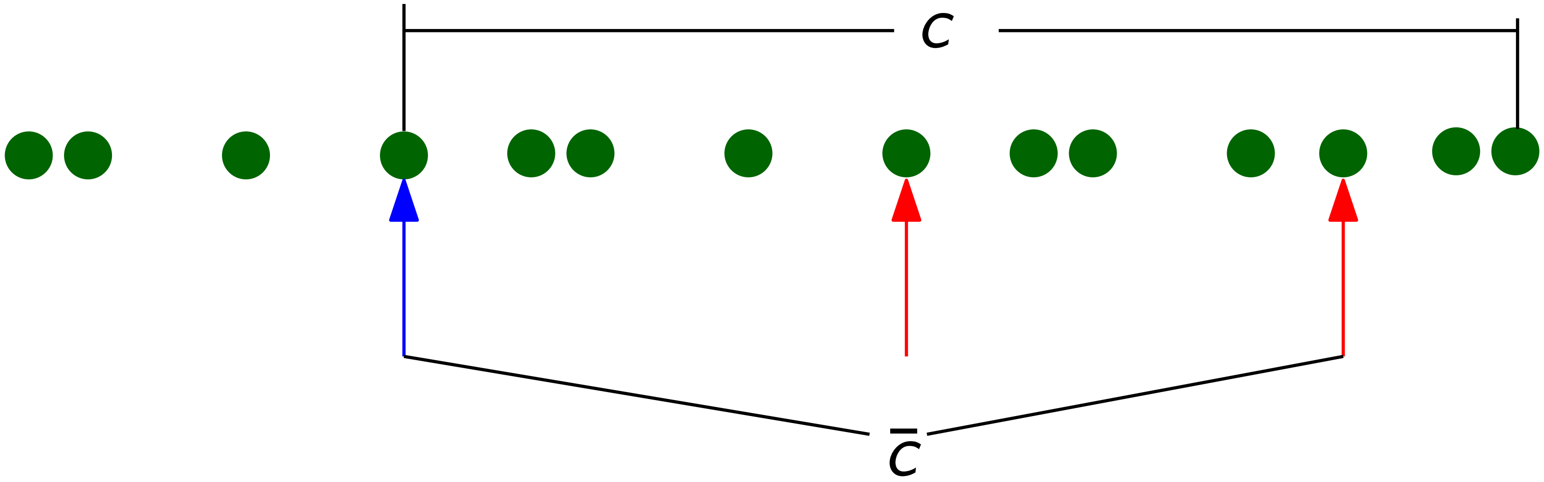

Coordinator only know $\bar{c}$

$$\hat{f}_i = \begin{cases} c - 1 + 1/p, & \text{if } c > 0; \\ 0, & \text{else.} \end{cases}$$

Frequent Items: analysis

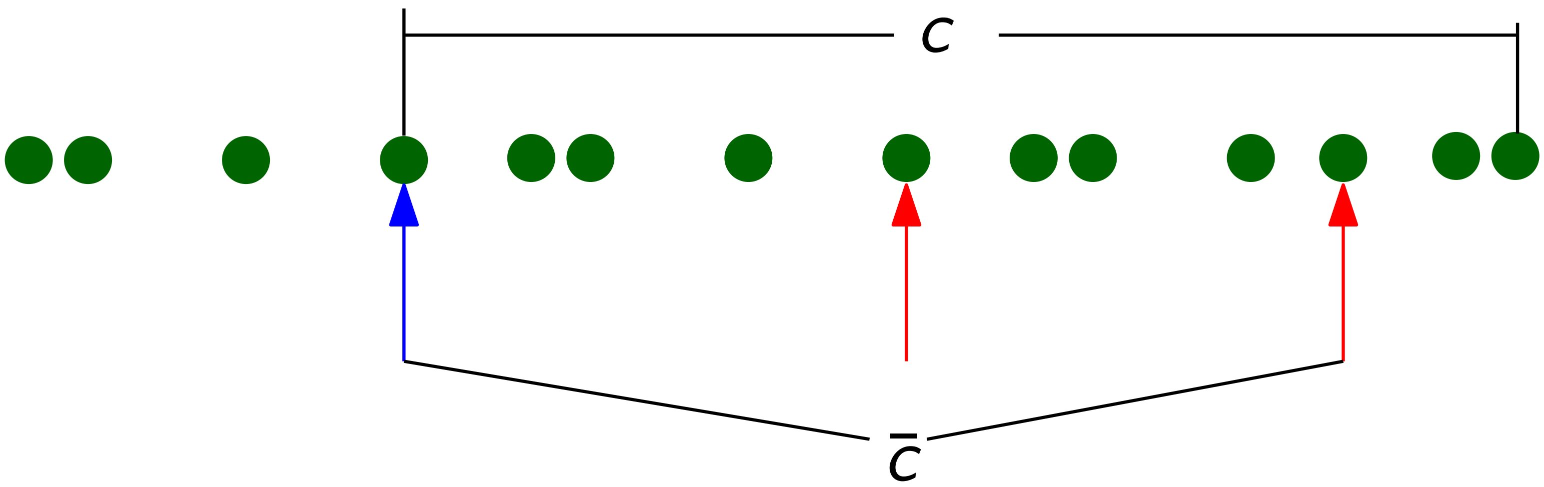

Estimate c by $\bar{c}$

$$\hat{c} = \begin{cases} \bar{c} - 1 + 1/p, & \text{if } \bar{c} > 0; \\ 0, & \text{else.} \end{cases}$$

Frequent Items: analysis

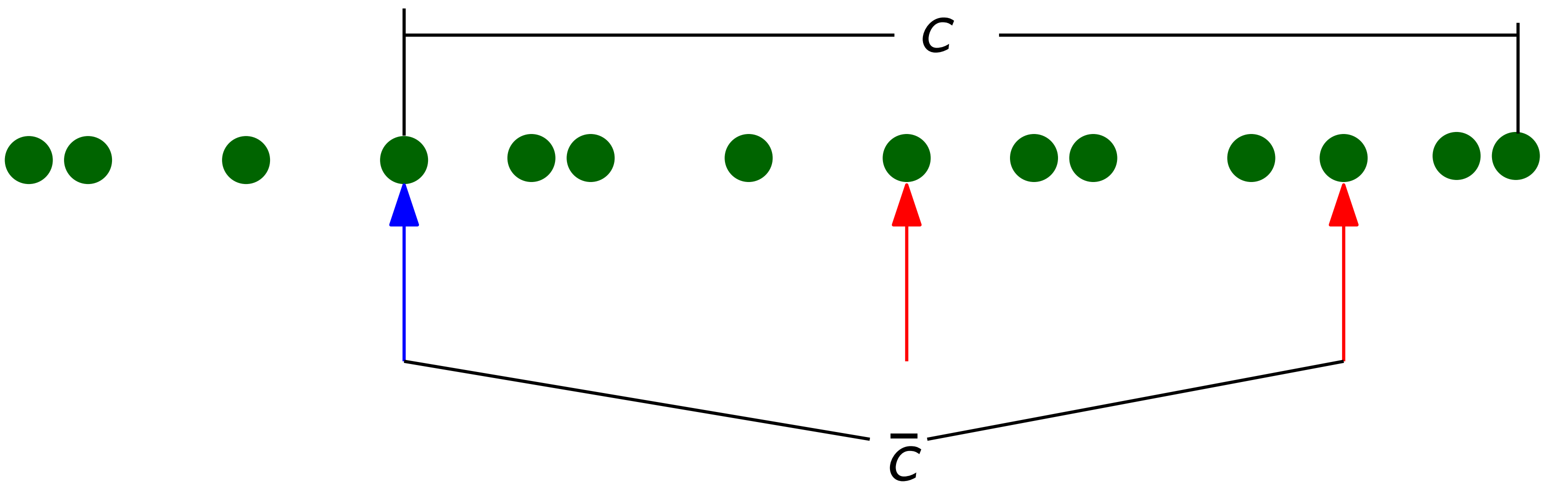

Combined

$$\hat{f}_i = \begin{cases} \bar{c} - 2 + 2/p, & \text{if } \bar{c} > 1; \\ 1/p, & \text{if } \bar{c} = 1; \\ 0, & \text{else.} \end{cases}$$

Lower bound

- $\mathrm{E}[\hat{f}_i] = f_i$
- $\mathrm{Var}[\hat{f}_i] \leq 2/p^2$

Lower bound

- $\mathrm{E}[\hat{f}_i] = f_i$
- $\mathrm{Var}[\hat{f}_i] \leq 2/p^2$

  set $p = O(\frac{\sqrt{k}}{\varepsilon n})$

  space: $O(\sqrt{k}/\varepsilon)$

  space per site: $O(1/(\varepsilon\sqrt{k}))$

Lower bound

- Communication lower bound still hold
- Space lower bound

Lower bound

- Communication lower bound still hold
- Space lower bound

    constant space

    communication space tradeoff

Space lower bound

Theorem

Any randomized algorithm that solves the frequency tracking problem with communication C bits and uses M bits of space per site, we have $C \cdot M = \Omega(\log n/\varepsilon^2)$.

Space lower bound

Theorem

Any randomized algorithm that solves the frequency tracking problem with communication C bits and uses M bits of space per site, we have $C \cdot M = \Omega(\log n/\varepsilon^2)$.

Communication cost: $O(\sqrt{k}/\varepsilon \cdot \log n)$ bits

Space per site: $\Omega(1/(\varepsilon\sqrt{k}))$ bits

Space lower bound

Theorem

The k-party communication complexity for the one-shot frequency estimation problem is $\Omega(\sqrt{k}/\varepsilon)$ bits.

[Woodruff, Zhang, STOC'12]

Space lower bound

Theorem

The k-party communication complexity for the one-shot frequency estimation problem is $\Omega(\sqrt{k}/\varepsilon)$ bits.

Direct-Sum theorem

Solve ℓ instances of the frequency estimation problem simultaneously needs $\Omega(\ell \cdot \sqrt{k}/\varepsilon)$ bits of communication.

[Woodruff, Zhang, STOC'12]

Space lower bound

Proof sketch

Let $\mathcal{A}$ be a k-party tracking algorithm with communication C and space M

Use $\mathcal{A}$ to solve tk-party one-shot problem.

Space lower bound

Proof sketch

Let $\mathcal{A}$ be a k-party tracking algorithm with communication C and space M

Use $\mathcal{A}$ to solve tk-party one-shot problem.

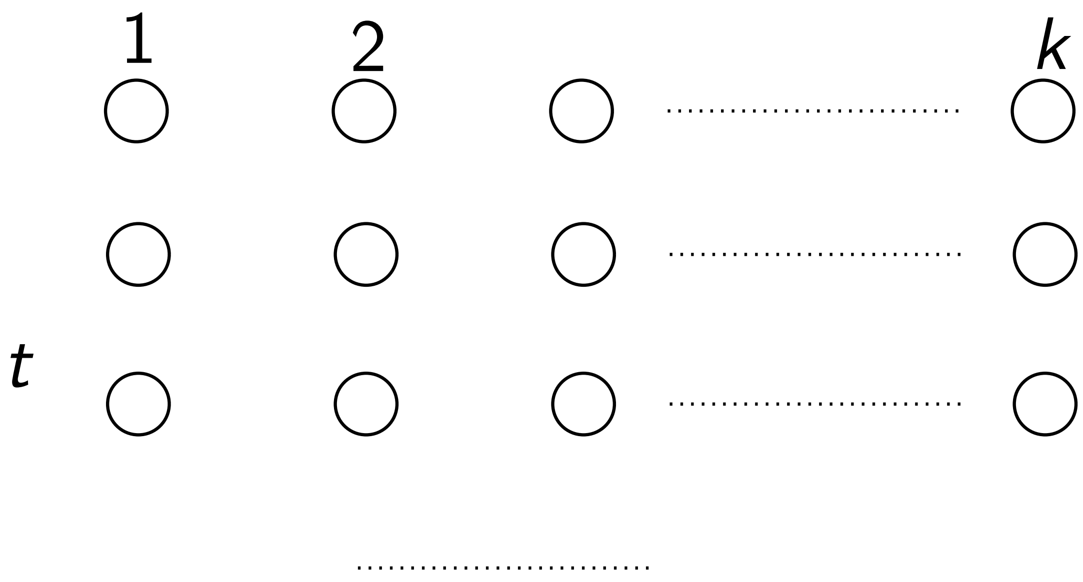

Space lower bound

Proof sketch

Let $\mathcal{A}$ be a k-party tracking algorithm with communication C and space M

Use $\mathcal{A}$ to solve tk-party one-shot problem.

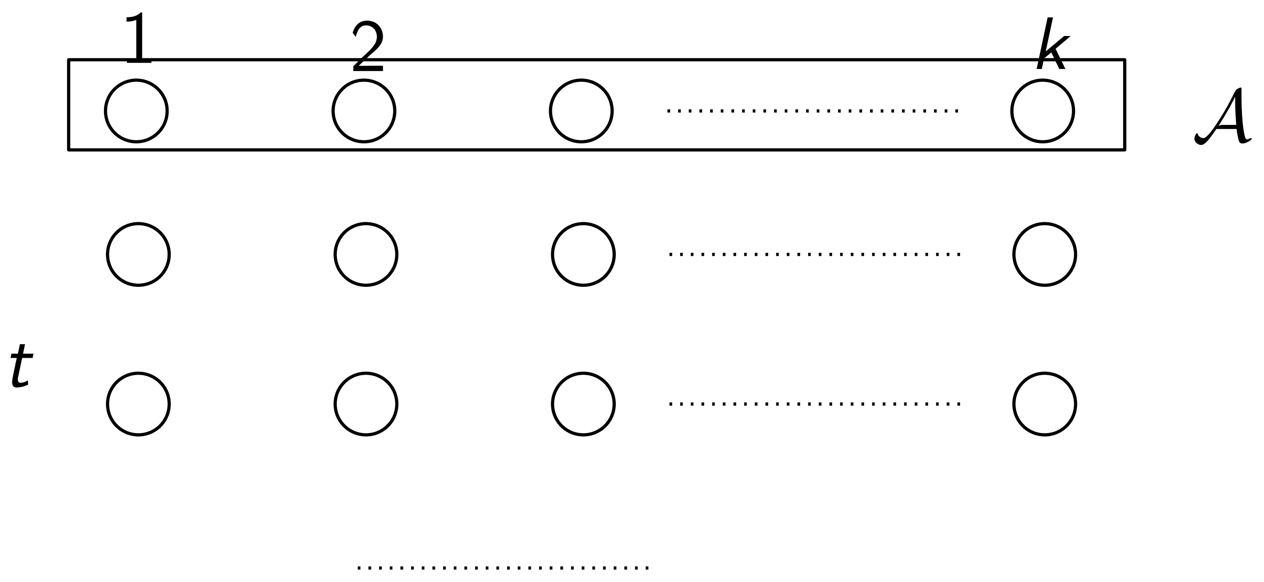

Space lower bound

Proof sketch

Let $\mathcal{A}$ be a k-party tracking algorithm with communication C and space M

Use $\mathcal{A}$ to solve tk-party one-shot problem.

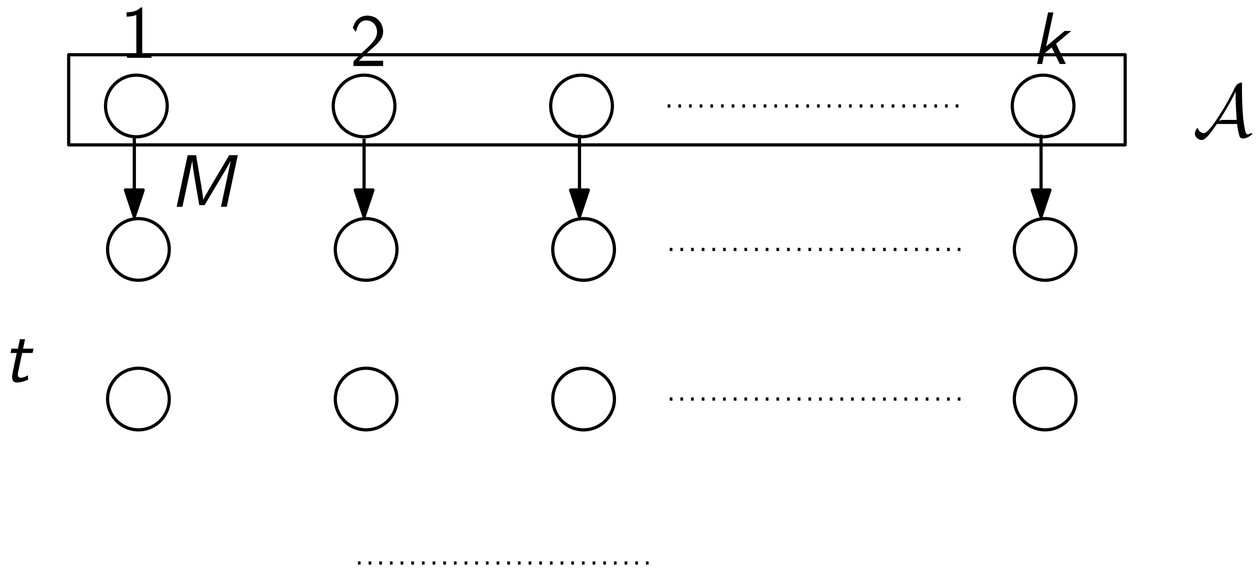

Space lower bound

Proof sketch

Let $\mathcal{A}$ be a k-party tracking algorithm with communication C and space M

Use $\mathcal{A}$ to solve tk-party one-shot problem.

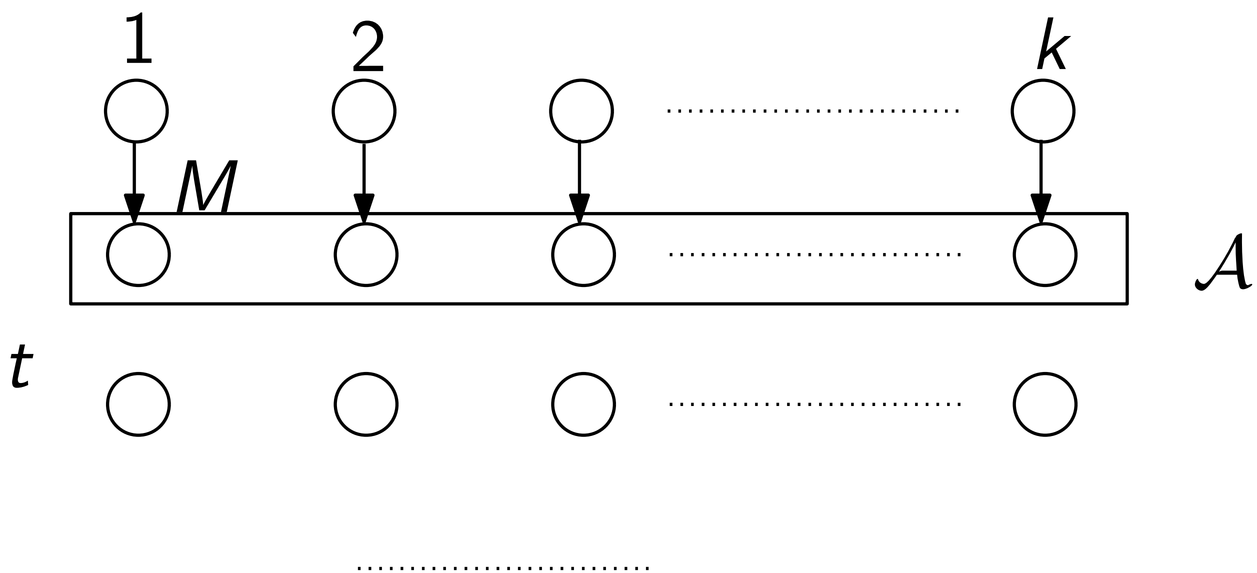

Space lower bound

Proof sketch

Let $\mathcal{A}$ be a k-party tracking algorithm with communication C and space M

Use $\mathcal{A}$ to solve tk-party one-shot problem.

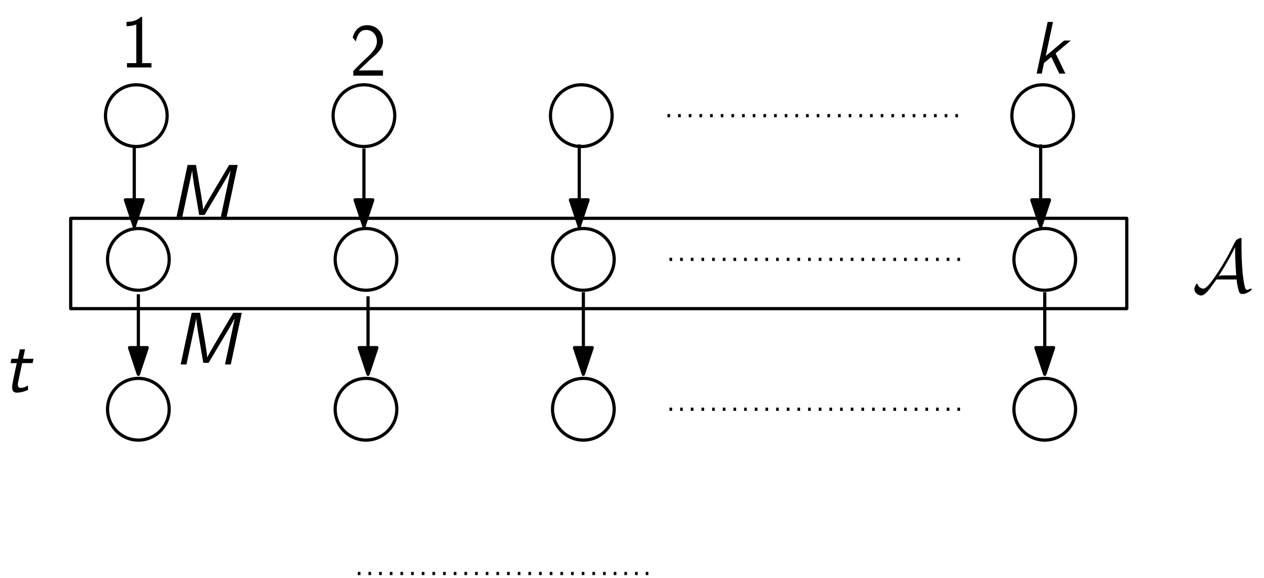

Space lower bound

Proof sketch

Let $\mathcal{A}$ be a k-party tracking algorithm with communication C and space M

Use $\mathcal{A}$ to solve tk-party one-shot problem.

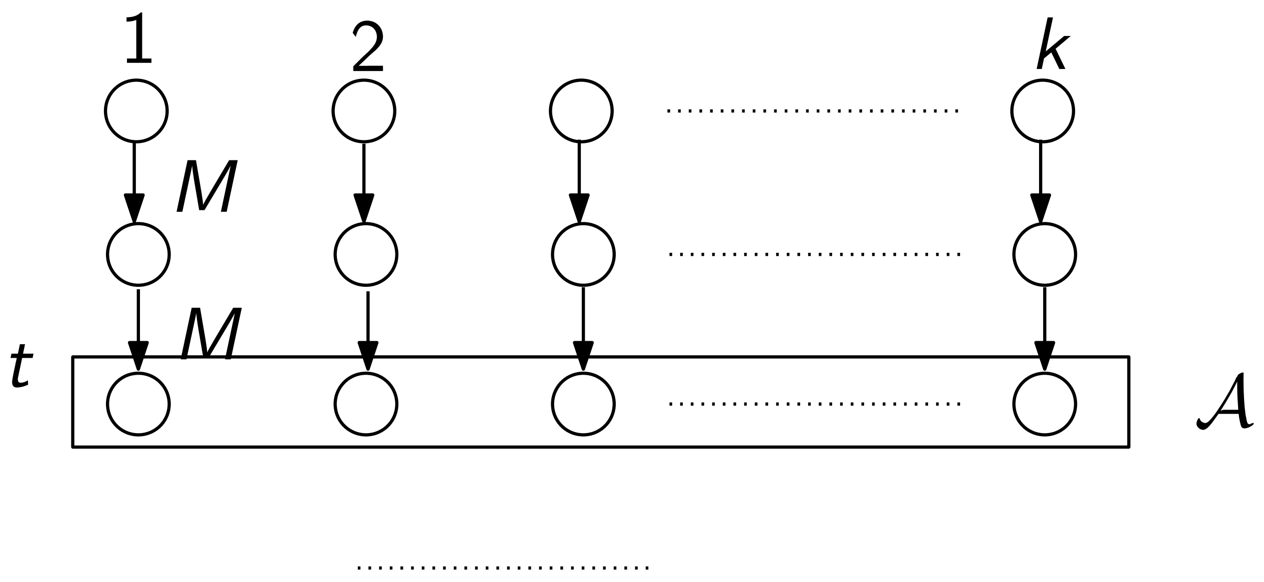

Space lower bound

Proof sketch

Let $\mathcal{A}$ be a k-party tracking algorithm with communication C and space M

Use $\mathcal{A}$ to solve tk-party one-shot problem.

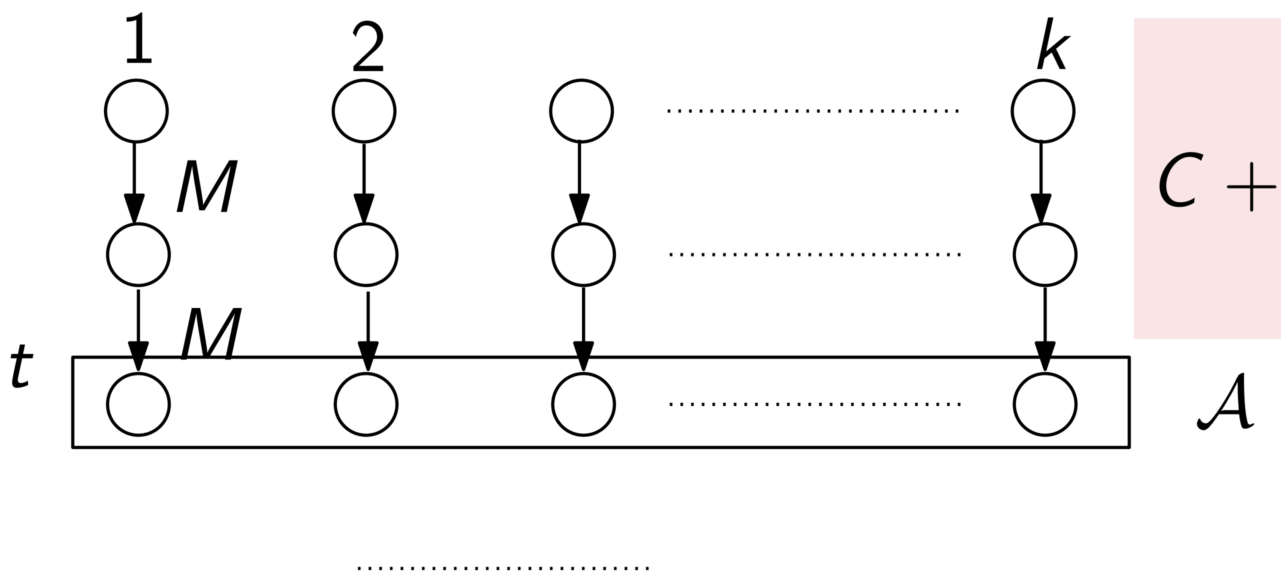

$$C + M \cdot tk \geq \Omega(\frac{\sqrt{kt}}{\varepsilon})$$

Approximate rank

value-to-rank queries

10 $r(10) = 7$

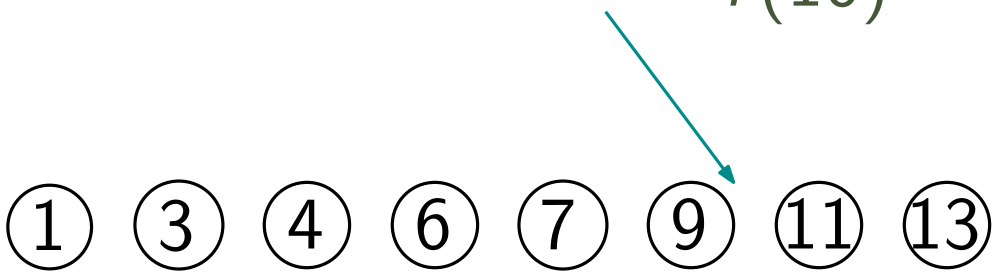

Return any number within $r(10) \pm \varepsilon n$

Approximate rank

n/k

Approximate rank

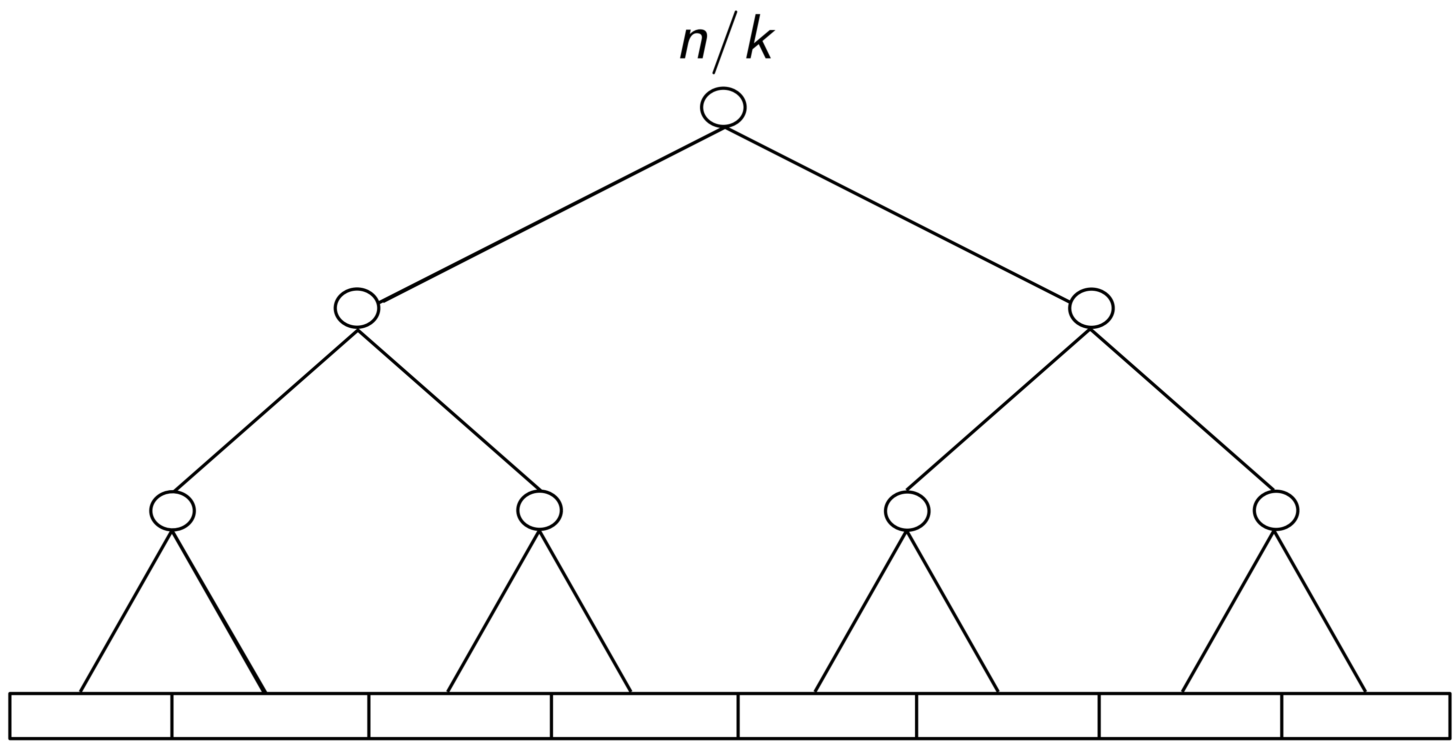

Block size $\frac{\varepsilon n}{\sqrt{k}}$

Approximate rank

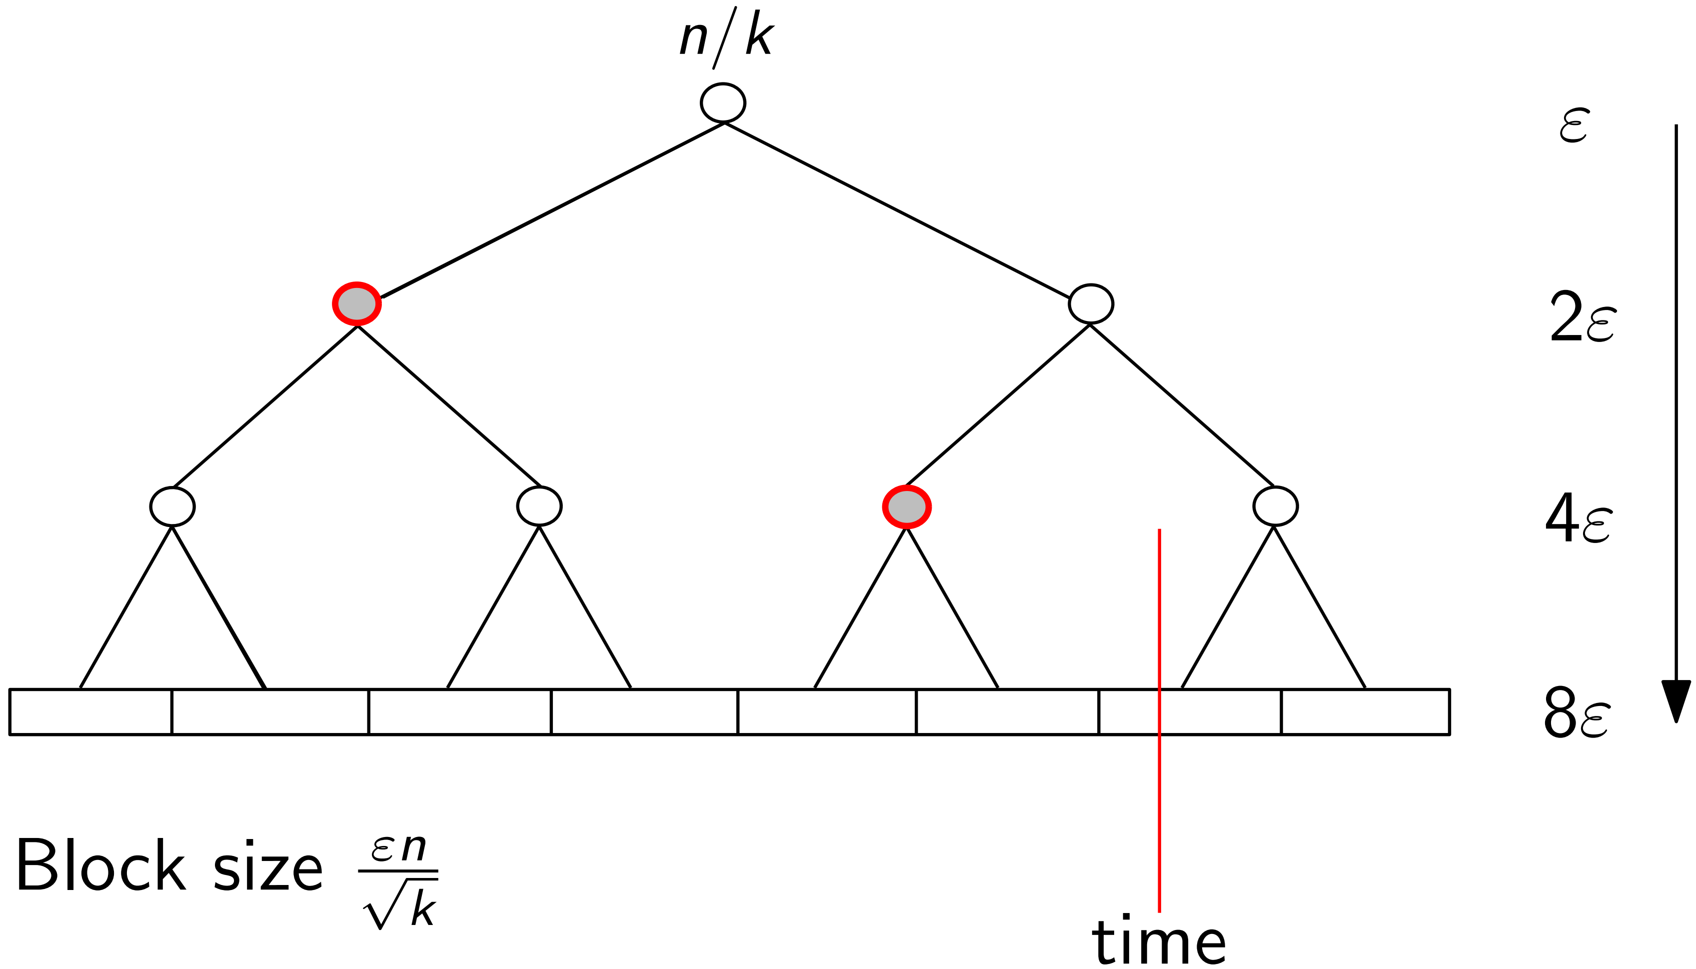

Block size $\frac{\varepsilon n}{\sqrt{k}}$

Each node will run a streaming algorithm

A node is full, send the sketch

Approximate rank

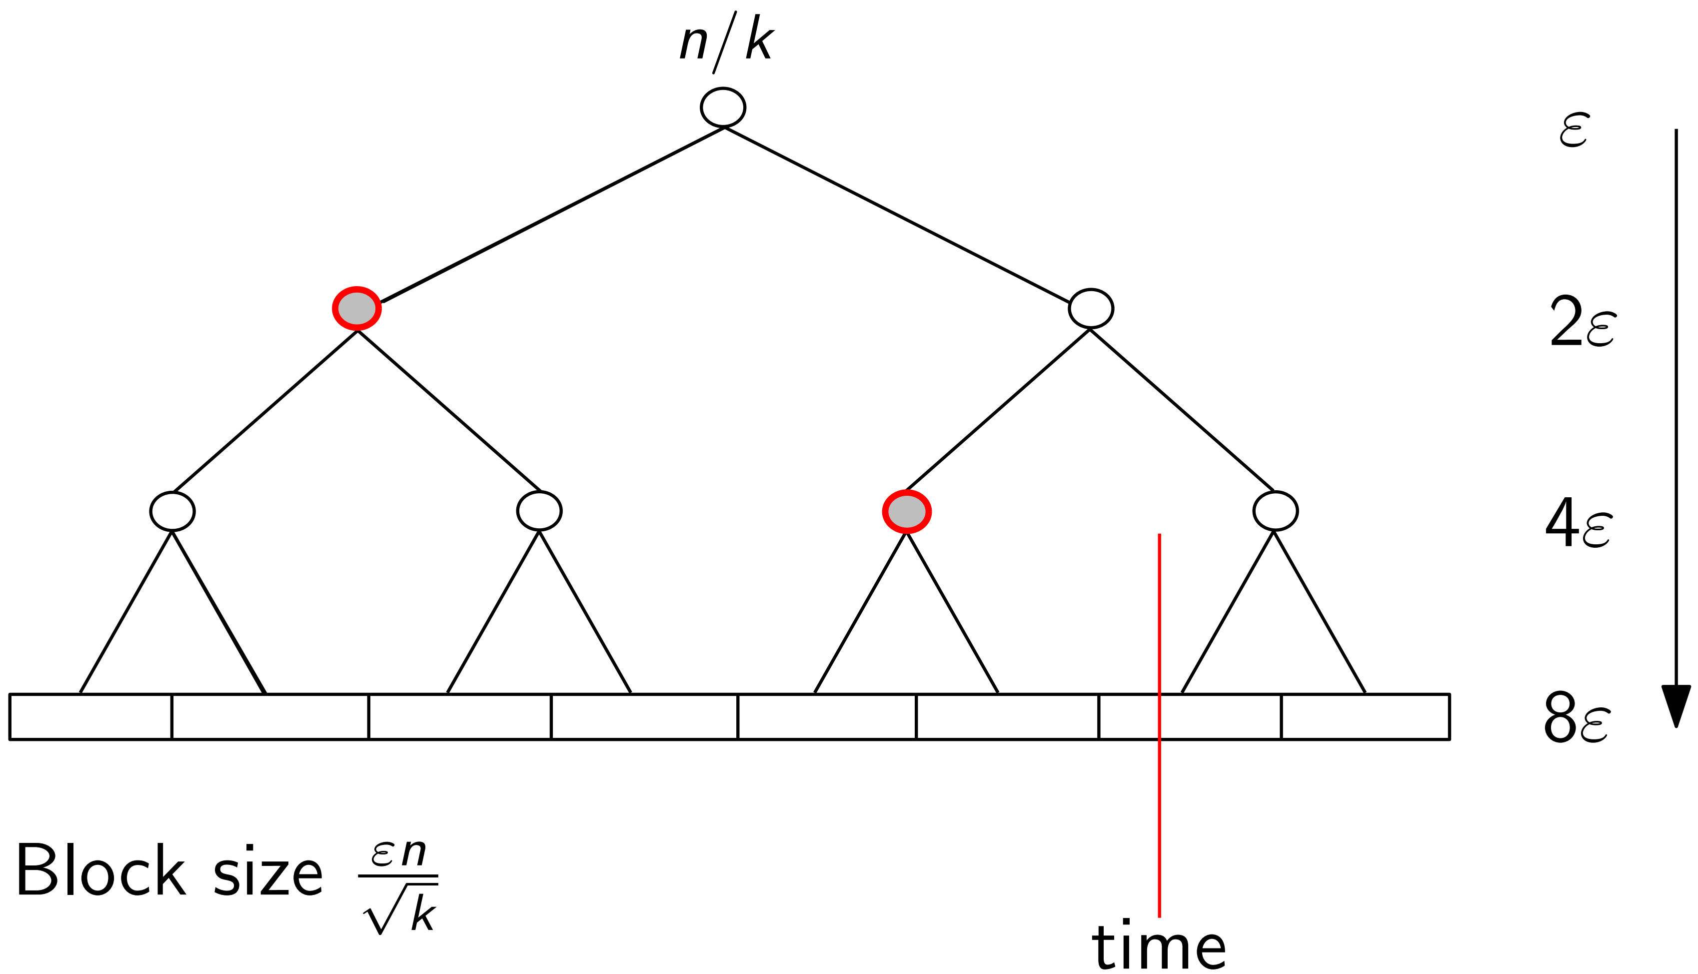

Block size $\frac{\varepsilon n}{\sqrt{k}}$

Analysis omitted

Approximate rank

Communication: $O(\frac{k}{\varepsilon} \cdot \log n \log^2 \frac{1}{\varepsilon})$
Space per site: $O(\frac{1}{\varepsilon} \cdot \log n)$

[Yi, Zhang, PODS'09]

Communication: $O(\frac{\sqrt{k}}{\varepsilon} \cdot \log n \log^{1.5} \frac{1}{\varepsilon})$
Space per site: $O(\frac{1}{\varepsilon\sqrt{k}}) \cdot \log^2 \frac{1}{\varepsilon}$

New

Conclusion

- Improve the communcation cost for 3 important problem by $O(\sqrt{k})$

- Improve the space by $O(\sqrt{k})$

- Prove matching randomized lower bound for communication and space

Thank you!