Tight Bounds for Distributed Streaming

(a.k.a., Distributed Functional Monitoring)

David Woodruff
IBM Research Almaden

Qin Zhang
MADALGO, Aarhus Univ.

STOC'12
May 22, 2012

The distributed streaming model

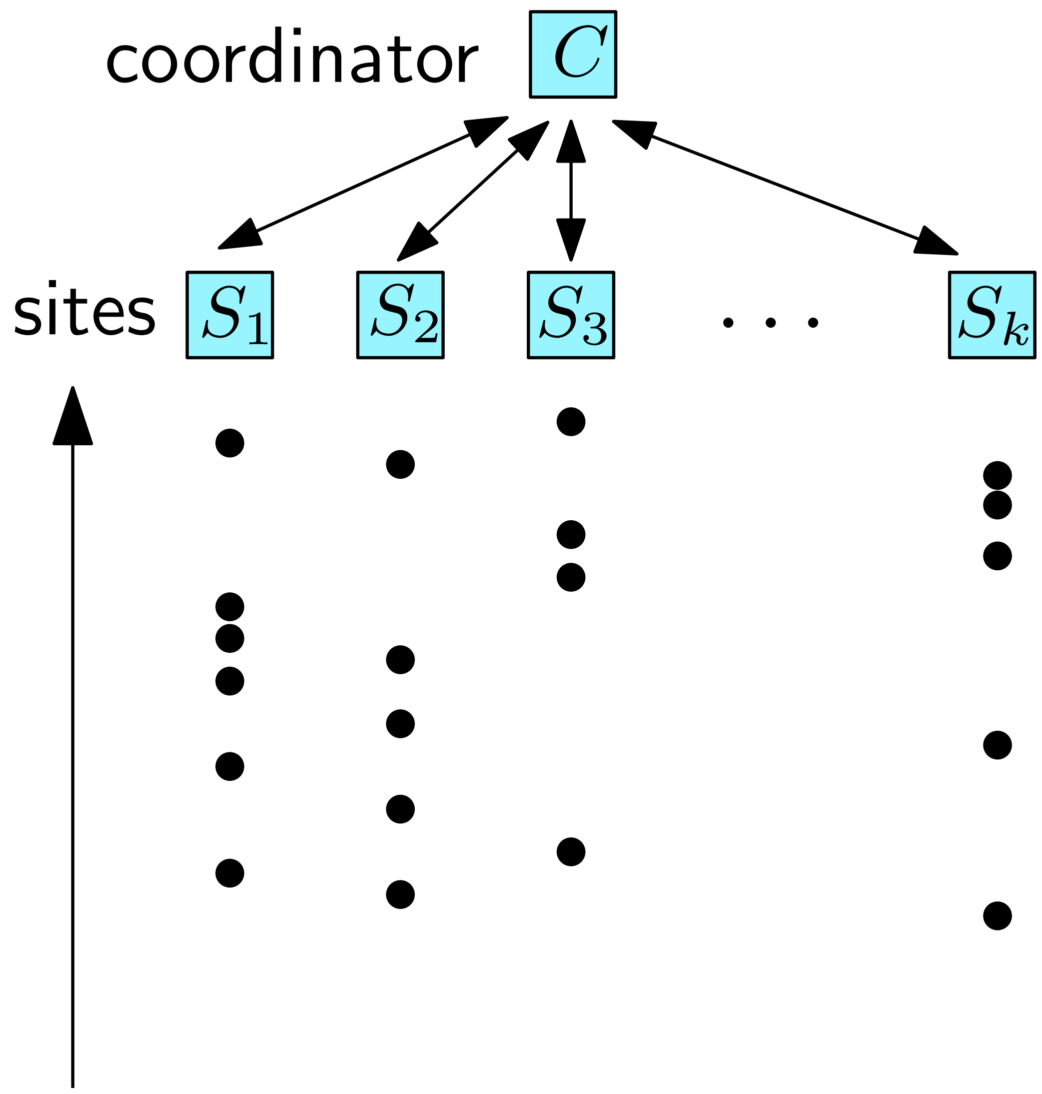

The coordinator needs to maintain some function defined on the union of k streams at any time

The distributed streaming model

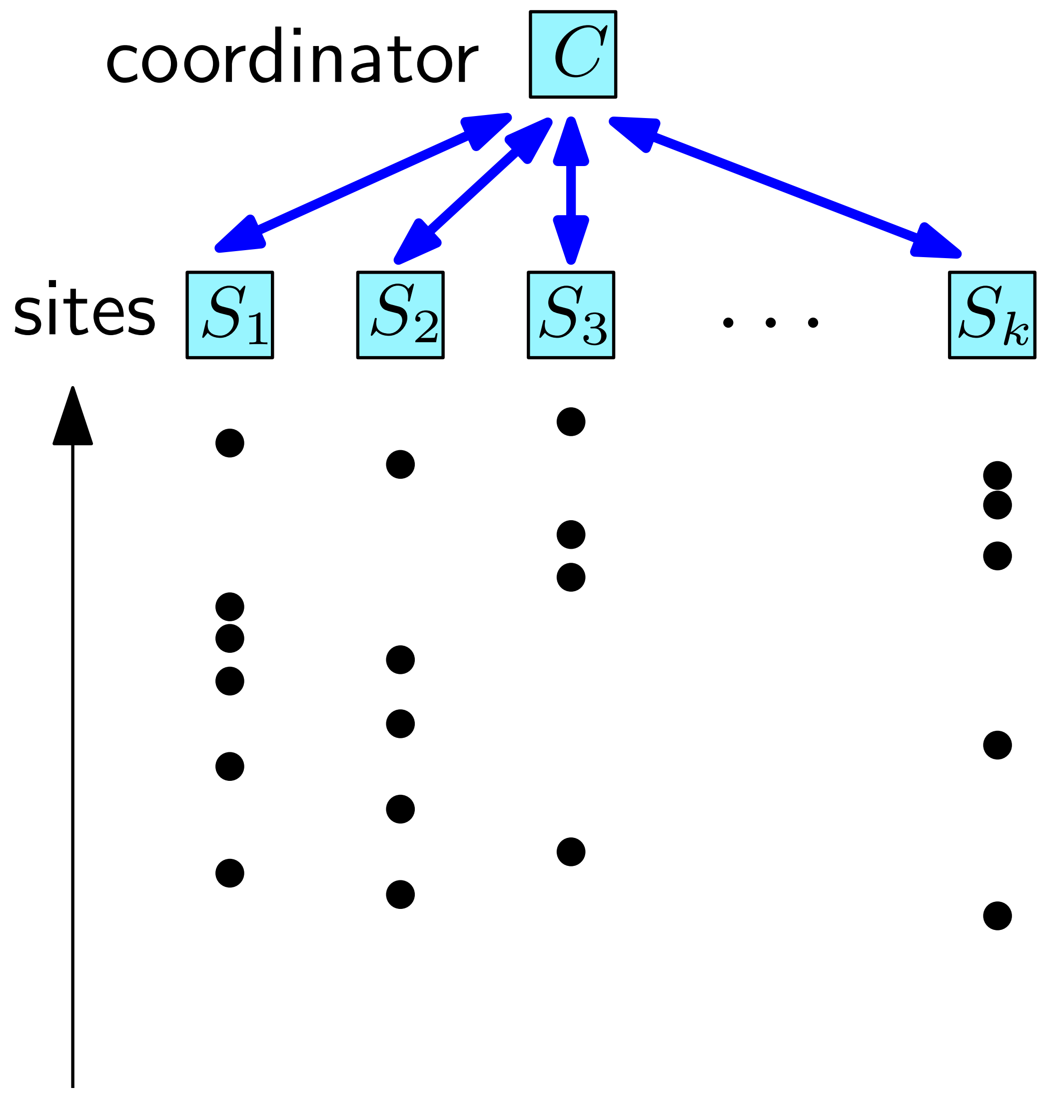

The coordinator needs to maintain some function defined on the union of k streams at any time

Goal: minimize total bits of communication

The distributed streaming model

Q: What's the total # items?

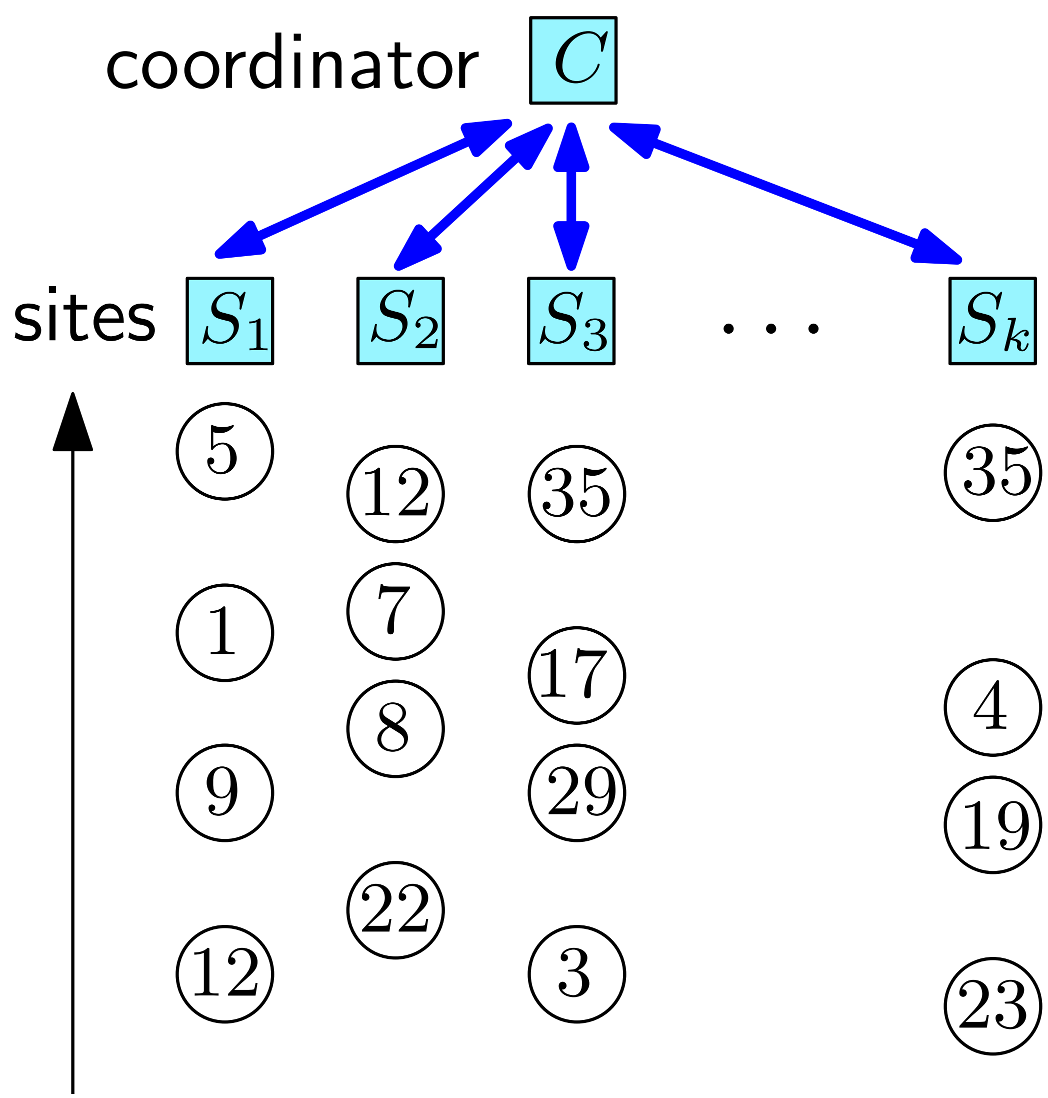

The coordinator needs to maintain some function defined on the union of k streams at any time

e.g., total # items

Goal: minimize total bits of communication

The distributed streaming model

Q: What's the total # items?

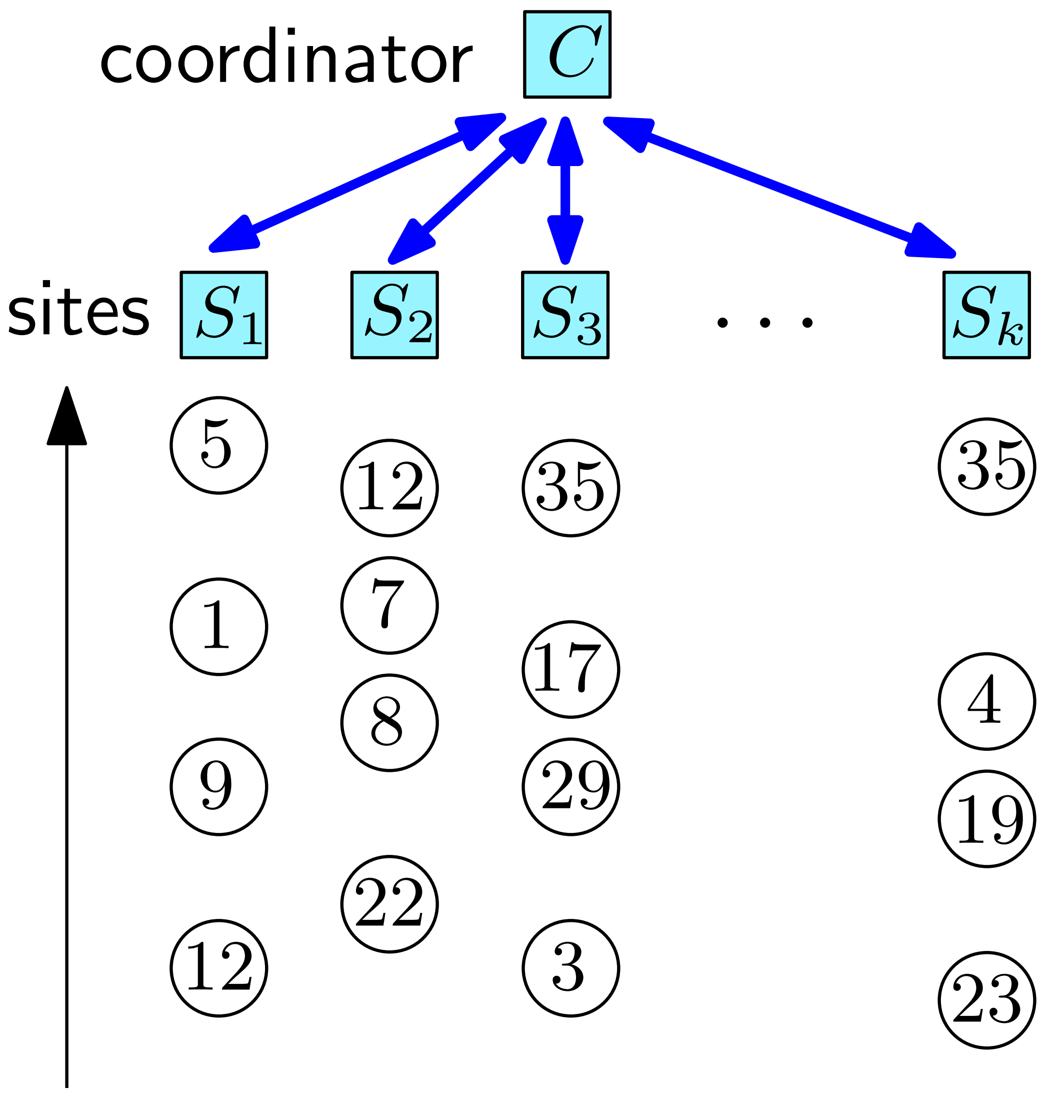

The coordinator needs to maintain some function defined on the union of k streams at any time

e.g., total # items

Almost always allow approximation.

Goal: minimize total bits of communication

Applied motivation: Distributed monitoring

Sensor networks

Sites: sensors

Streams:
e.g., environmental data

Concerns: energy

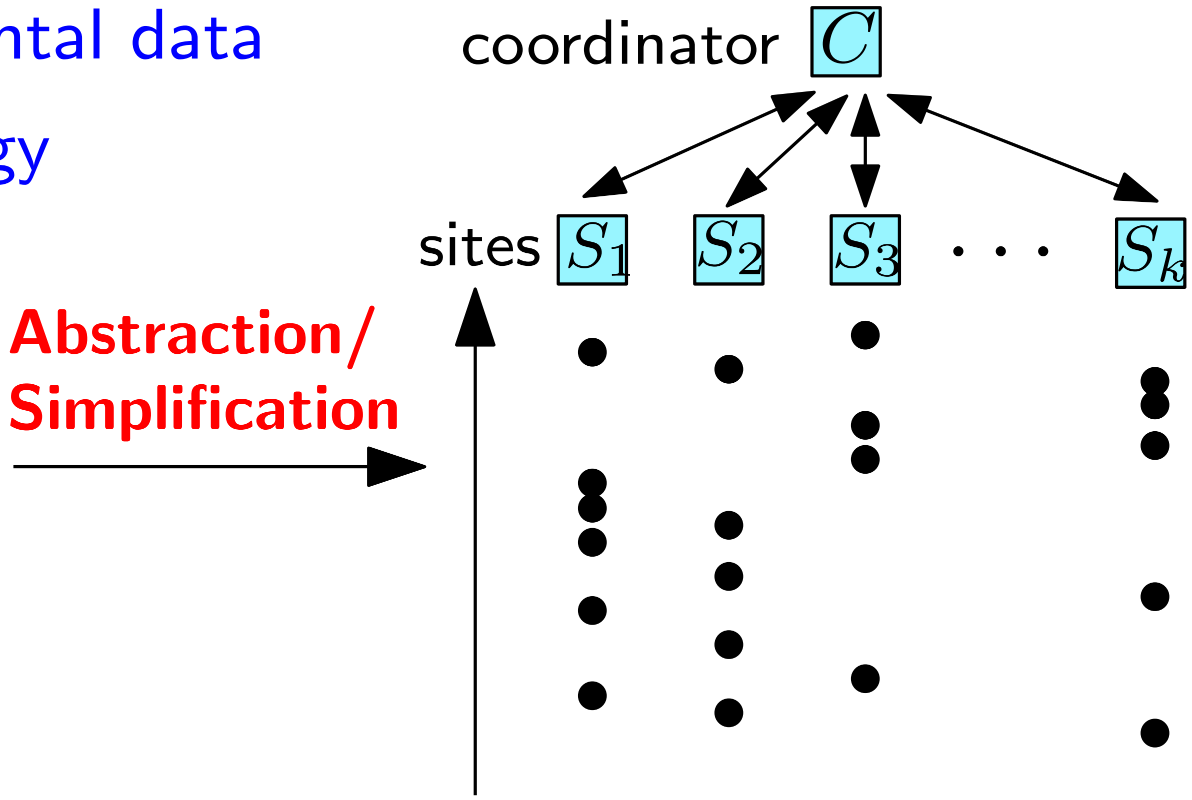

Applied motivation: Distributed monitoring

Sensor networks

Sites: sensors

Streams:
e.g., environmental data

Concerns: energy

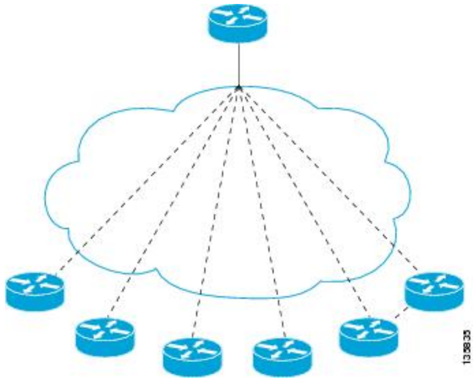

Network routers

Sites: routers

Streams:
e.g., IP addresses.

Concerns: bandwidth

Data in the cloud

Sites: machines

Streams:
e.g., queries/updates.

Concerns: bandwidth

Abstraction/ Simplification

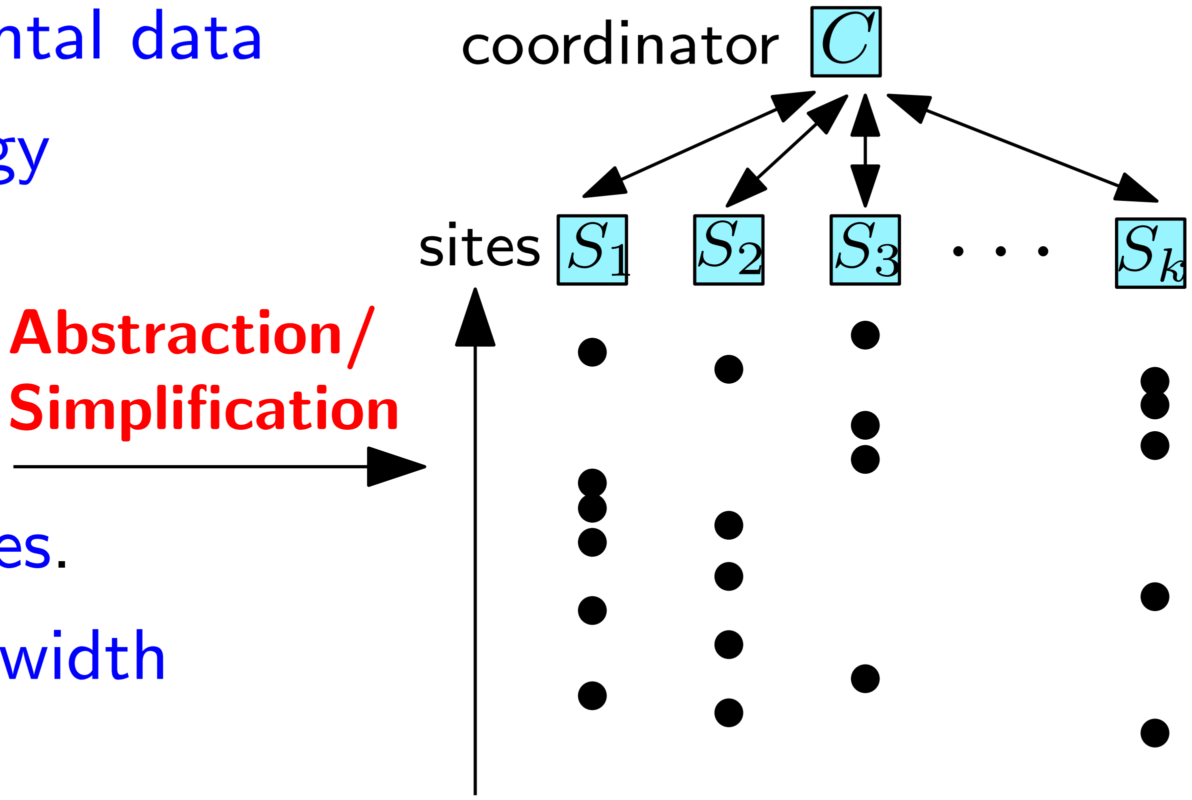

Problems

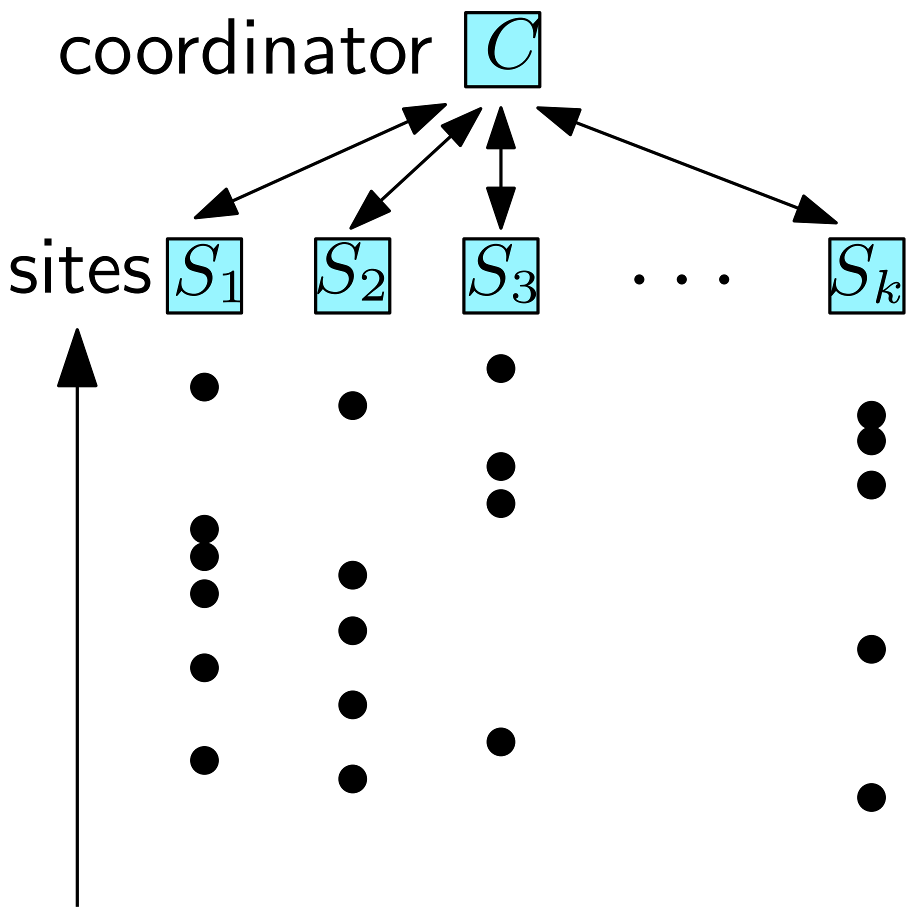

The Distributed Streaming Model

- Frequency moments

  $F_p = \sum_i f_i^p$

  f_i: freq. of element i

  In particular:

  F_0: #distinct elements

  F_2: size of self-join

- Heavy hitters
- Quantile
- Entropy
- ...

Well-studied problems in the data stream literature

Results

k: # sites; n: input size;
ε: approximation ratio

Problems	Upper Bound	Lower Bound
F_p $(p > 1)$	$\tilde{O}(k^{2p+1}n^{1-2/p}/\text{poly}(\varepsilon))$ [CMY, SODA '08] $\tilde{O}(k^{p-1}/\text{poly}(\varepsilon))$ [This paper]	$\Omega(k)$ [CMY, SODA '08] $\tilde{\Omega}(k^{p-1}/\varepsilon^2)$ [This paper]
F_0	$\tilde{O}(k/\varepsilon^2)$ [CMY, SODA '08]	$\Omega(k)$ [CMY, SODA '08] $\tilde{\Omega}(k/\varepsilon^2)$ [This paper]
Heavy-hitter/ Quantile	$\tilde{O}(\min\{\sqrt{k}/\varepsilon, 1/\varepsilon^2\})$ [HYZ, PODS '12]	$\tilde{\Omega}(\max\{\sqrt{k}/\varepsilon, 1/\varepsilon^2\})$ [HYZ, PODS '12]; [This paper] static case
Entropy	$\tilde{O}(k/\varepsilon^3)$ [ABC, ICALP '09]	$\Omega(1/\sqrt{\varepsilon})$ [ABC, ICALP '09] $\tilde{\Omega}(k/\varepsilon^2)$ [This paper]

– Improve LBs for all problems, and the UB for F_p $(p > 1)$.

Results

k: # sites; n: input size;
ε: approximation ratio

Problems	Upper Bound	Lower Bound
F_p $(p > 1)$	$\tilde{O}(k^{2p+1} n^{1-2/p}/\mathrm{poly}(\varepsilon))$ [CMY, SODA '08] $\tilde{O}(k^{p-1}/\mathrm{poly}(\varepsilon))$ [This paper]	$\Omega(k)$ [CMY, SODA '08] $\tilde{\Omega}(k^{p-1}/\varepsilon^2)$ [This paper]
F_0	$\tilde{O}(k/\varepsilon^2)$ [CMY, SODA '08]	$\Omega(k)$ [CMY, SODA '08] $\tilde{\Omega}(k/\varepsilon^2)$ [This paper]
Heavy-hitter/ Quantile	$\tilde{O}(\min\{\sqrt{k}/\varepsilon, 1/\varepsilon^2\})$ [HYZ, PODS '12]	$\tilde{\Omega}(\max\{\sqrt{k}/\varepsilon, 1/\varepsilon^2\})$ [HYZ, PODS '12]; [This paper] static case
Entropy	$\tilde{O}(k/\varepsilon^3)$ [ABC, ICALP '09]	$\Omega(1/\sqrt{\varepsilon})$ [ABC, ICALP '09] $\tilde{\Omega}(k/\varepsilon^2)$ [This paper]

– Improve LBs for all problems, and the UB for F_p $(p > 1)$.

– Our **LBs** even **hold in the static case**.
Static LBs (almost) **match continuous UBs**.

By-products

Implications for problems in the data stream model

- First lower bound for F_0 without Gap-Hamming
- Improve $\Omega(n^{1-2/p}/\varepsilon^{2/p}t)$ bound for estimating F_p $(p \geq 2)$ in a stream using t passes to $\Omega(n^{1-2/p}/\varepsilon^{4/p}t)$.

  First LB that agrees with the UB for F_2 $(p = 2)$, for any constant t.

The multiparty NIH communication model

– A model for (static) lower bounds

$x_1 = 010011$ $x_2 = 111011$

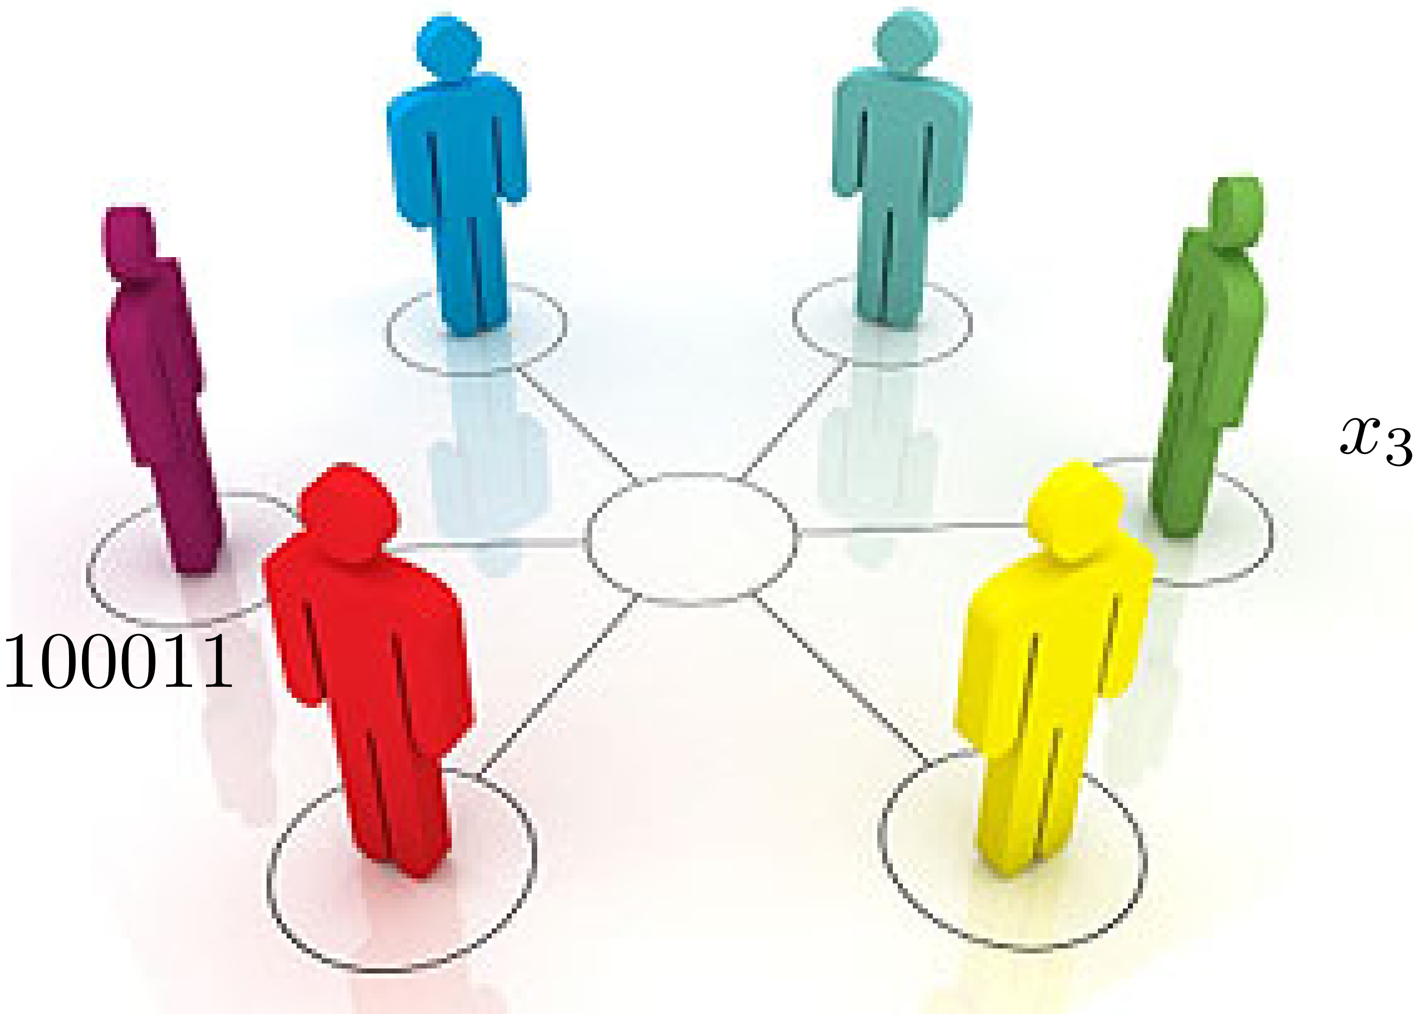

$x_3 = 111111$

$x_k = 100011$

They want to jointly compute $f(x_1, x_2, \ldots, x_k)$

Goal: minimize total bits of communication

The multiparty NIH communication model

– A model for (static) lower bounds

Blackboard: One speaks, everyone else hears.

Message passing: If x_1 talks to x_2, others cannot hear.

$x_1 = 010011$ $x_2 = 111011$

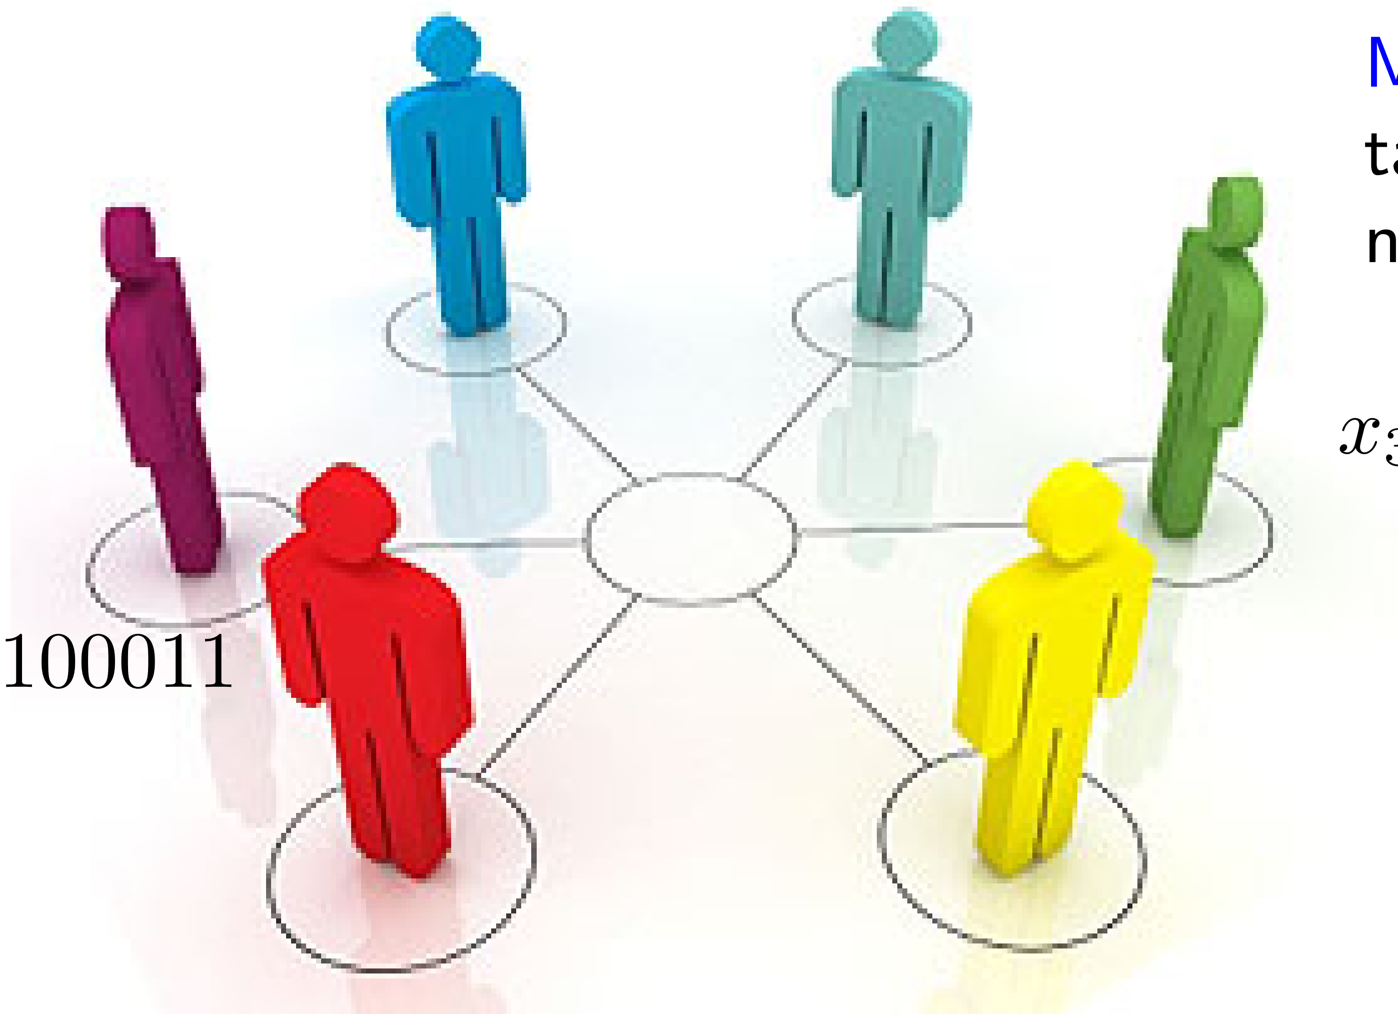

$x_3 = 111111$

$x_k = 100011$

They want to jointly compute $f(x_1, x_2, \ldots, x_k)$

Goal: minimize total bits of communication

The multiparty NIH communication model

– A model for (static) lower bounds

Blackboard: One speaks, everyone else hears.

Message passing: If x_1 talks to x_2, others cannot hear.

$x_1 = 010011$ $x_2 = 111011$

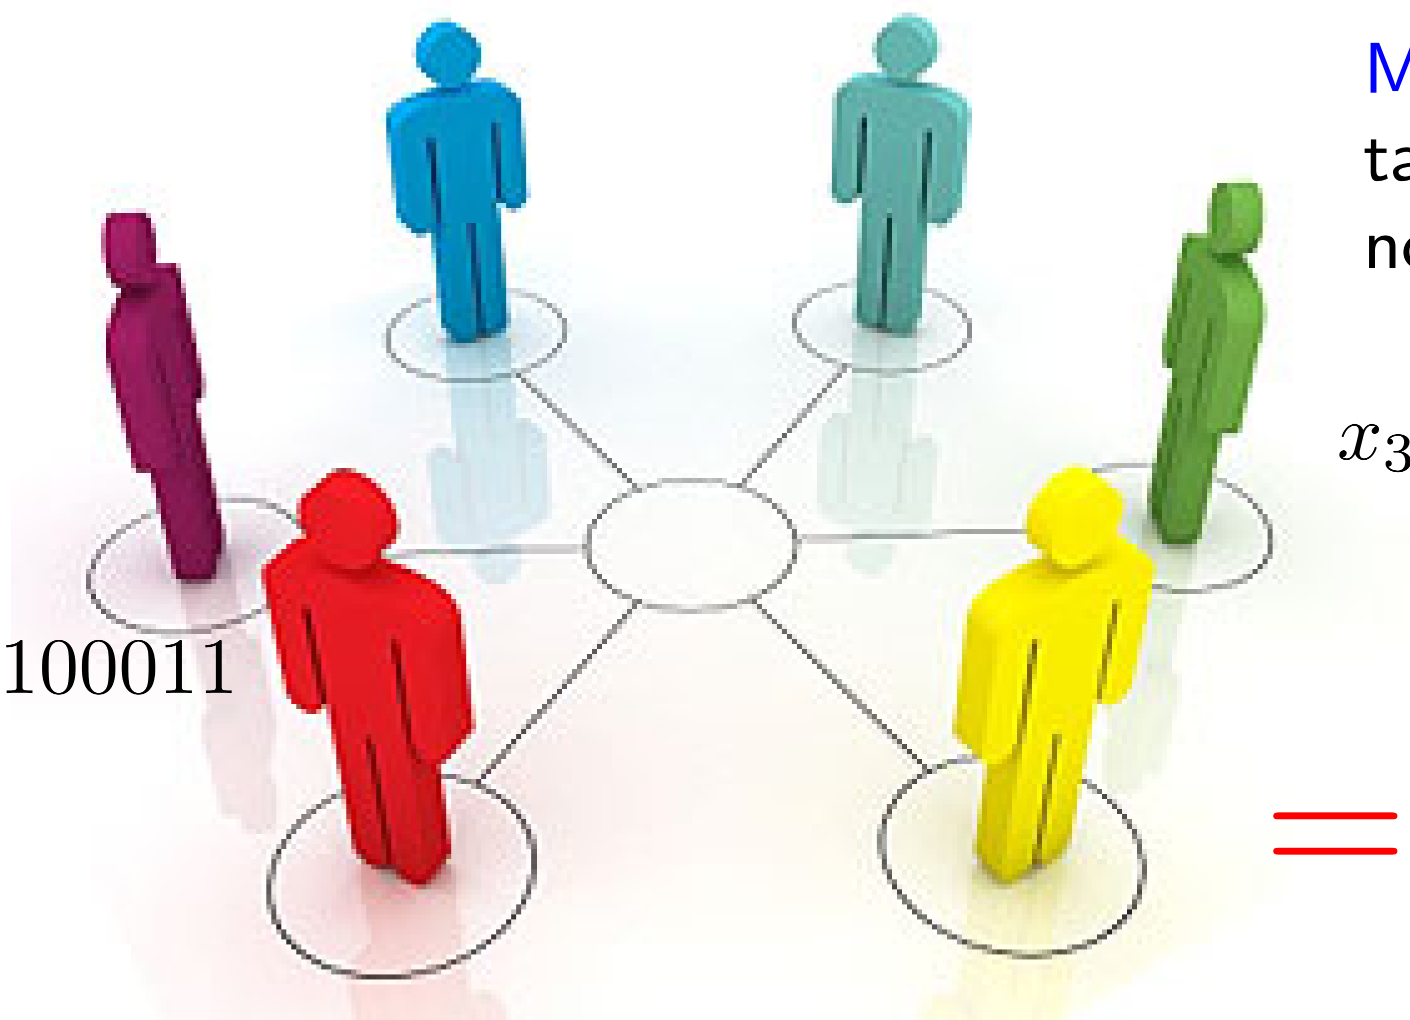

$x_3 = 111111$

$x_k = 100011$

$=$ static version of

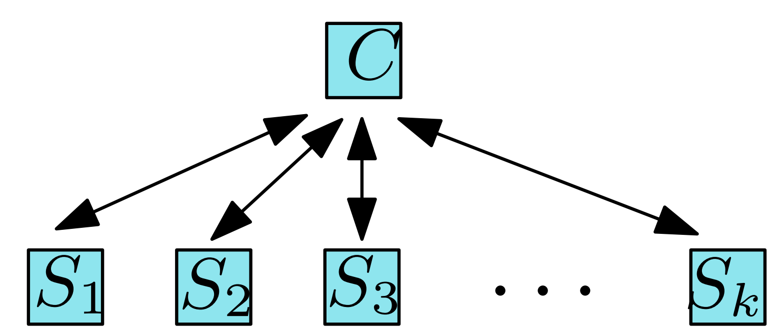

They want to jointly compute $f(x_1, x_2, \ldots, x_k)$

Goal: minimize total bits of communication

Previously, for multiparty NIH comm. model

- Well studied?

Previously, for multiparty NIH comm. model

- Well studied?

  YES and NO.

  Several papers in blackboard model.
  [Alon, Matias, Szegedy '96]
  [Bar-Yossef, Jayram, Kumar, Sivakumar '04]
  . . .

Previously, for multiparty NIH comm. model

- Well studied?

  YES and NO.

  Several papers in blackboard model.
  [Alon, Matias, Szegedy '96]
  [Bar-Yossef, Jayram, Kumar, Sivakumar '04]
  . . .

  Almost nothing in message-passing model.

Previously, for multiparty NIH comm. model

- Well studied?

  YES and NO.

  Several papers in blackboard model.
  [Alon, Matias, Szegedy '96]
  [Bar-Yossef, Jayram, Kumar, Sivakumar '04]
  ...

  Almost nothing in message-passing model.

- A technique called "symmetrization" is proposed [Phillips, Verbin, Zhang '12] which works for both variants.

  However, this technique cannot be applied to our problems, due to several inherent limitations.

Now our new technique:

Composition

Now our new technique:
Composition

A Lower bound for F_0 (distinct elements)

Will "cheat" constantly,
see our paper for the real proof

The F_0 problem

k sites each holds a set X_i $(i \in \{1, 2, \ldots, k\})$.

Goal: compute #distinct elements$(\cup_{i=1}^{k} X_i)$ up to a $(1+\varepsilon)$-approx.

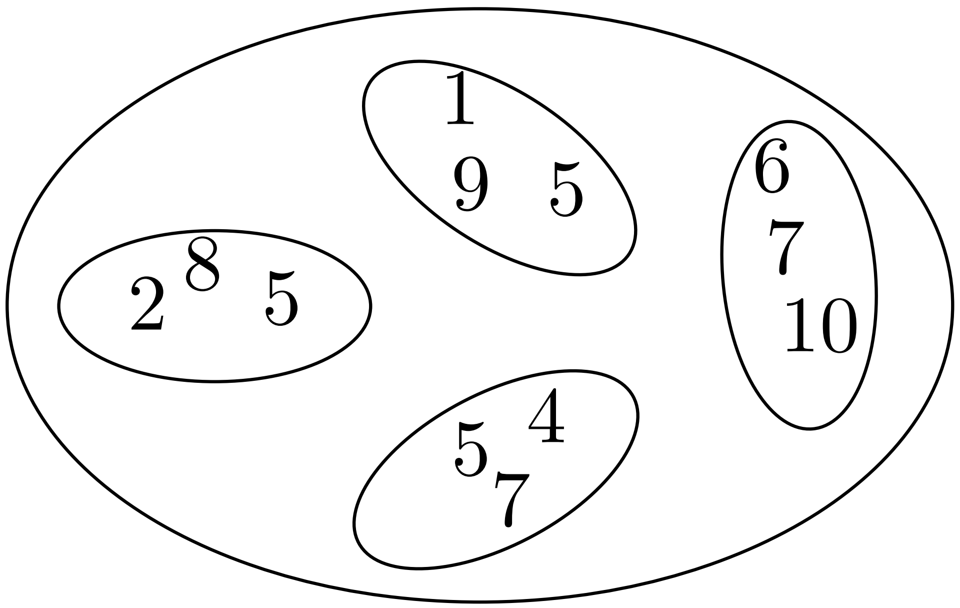

How many distinct items?

The F_0 problem

k sites each holds a set X_i $(i \in \{1, 2, \ldots, k\})$.

Goal: compute #distinct elements$(\cup_{i=1}^{k} X_i)$ up to a $(1+\varepsilon)$-approx.

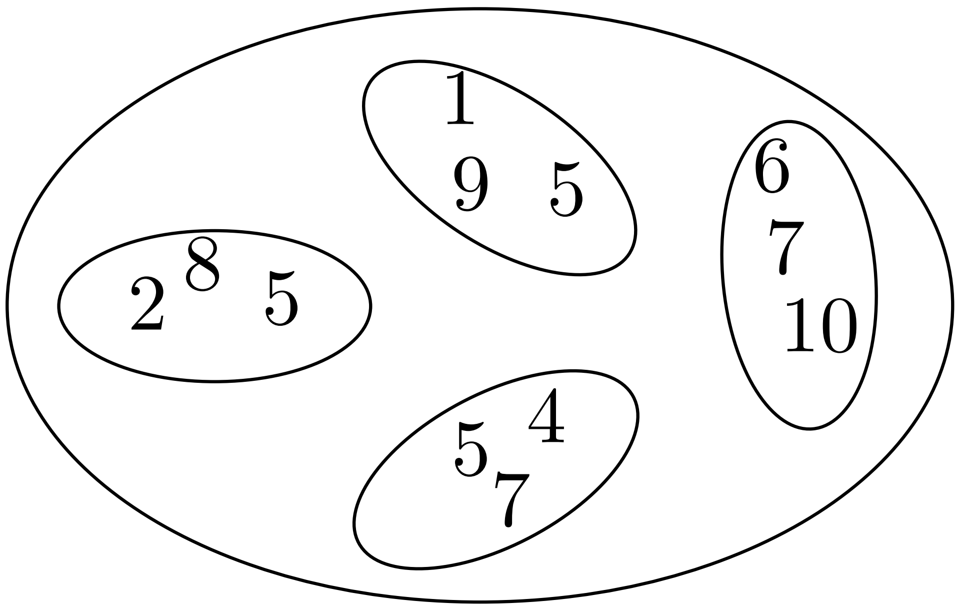

How many distinct items?

Current best UB: $\tilde{O}(k/\varepsilon^2)$
Holds in message-passing model

Previous LB: $\Omega(k)$

The F_0 problem

k sites each holds a set X_i $(i \in \{1, 2, \ldots, k\})$.

Goal: compute #distinct elements$(\cup_{i=1}^{k} X_i)$ up to a $(1+\varepsilon)$-approx.

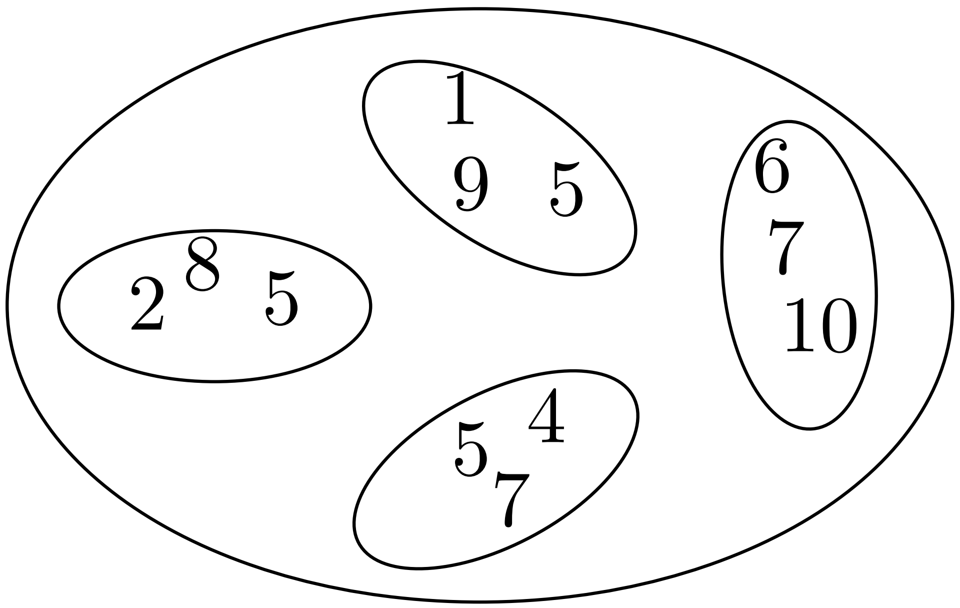

How many distinct items?

Current best UB: $\tilde{O}(k/\varepsilon^2)$
Holds in message-passing model

Previous LB: $\Omega(k)$

And $\Omega(1/\varepsilon^2)$
(reduction from Gap-Hamming)

The F_0 problem

k sites each holds a set X_i $(i \in \{1, 2, \ldots, k\})$.

Goal: compute #distinct elements$(\cup_{i=1}^{k} X_i)$ up to a $(1+\varepsilon)$-approx.

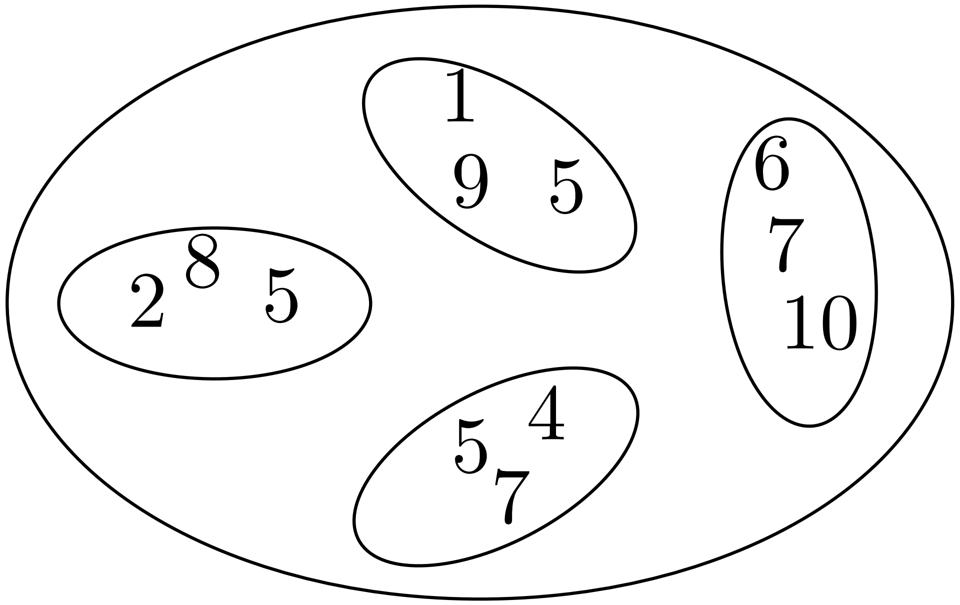

How many distinct items?

Current best UB: $\tilde{O}(k/\varepsilon^2)$
Holds in message-passing model

Previous LB: $\Omega(k)$

And $\Omega(1/\varepsilon^2)$
(reduction from Gap-Hamming)

Tight!

Our LB: $\Omega(k/\varepsilon^2)$.
Holds in message-passing model
Better UB in the blackboard model

F_0 hard input distribution

	1	2	•••	n
S_1	1	0	•••	0
S_2	0	1	•••	0
⋮			•••	
S_k	0	0	•••	0

F_0 hard input distribution

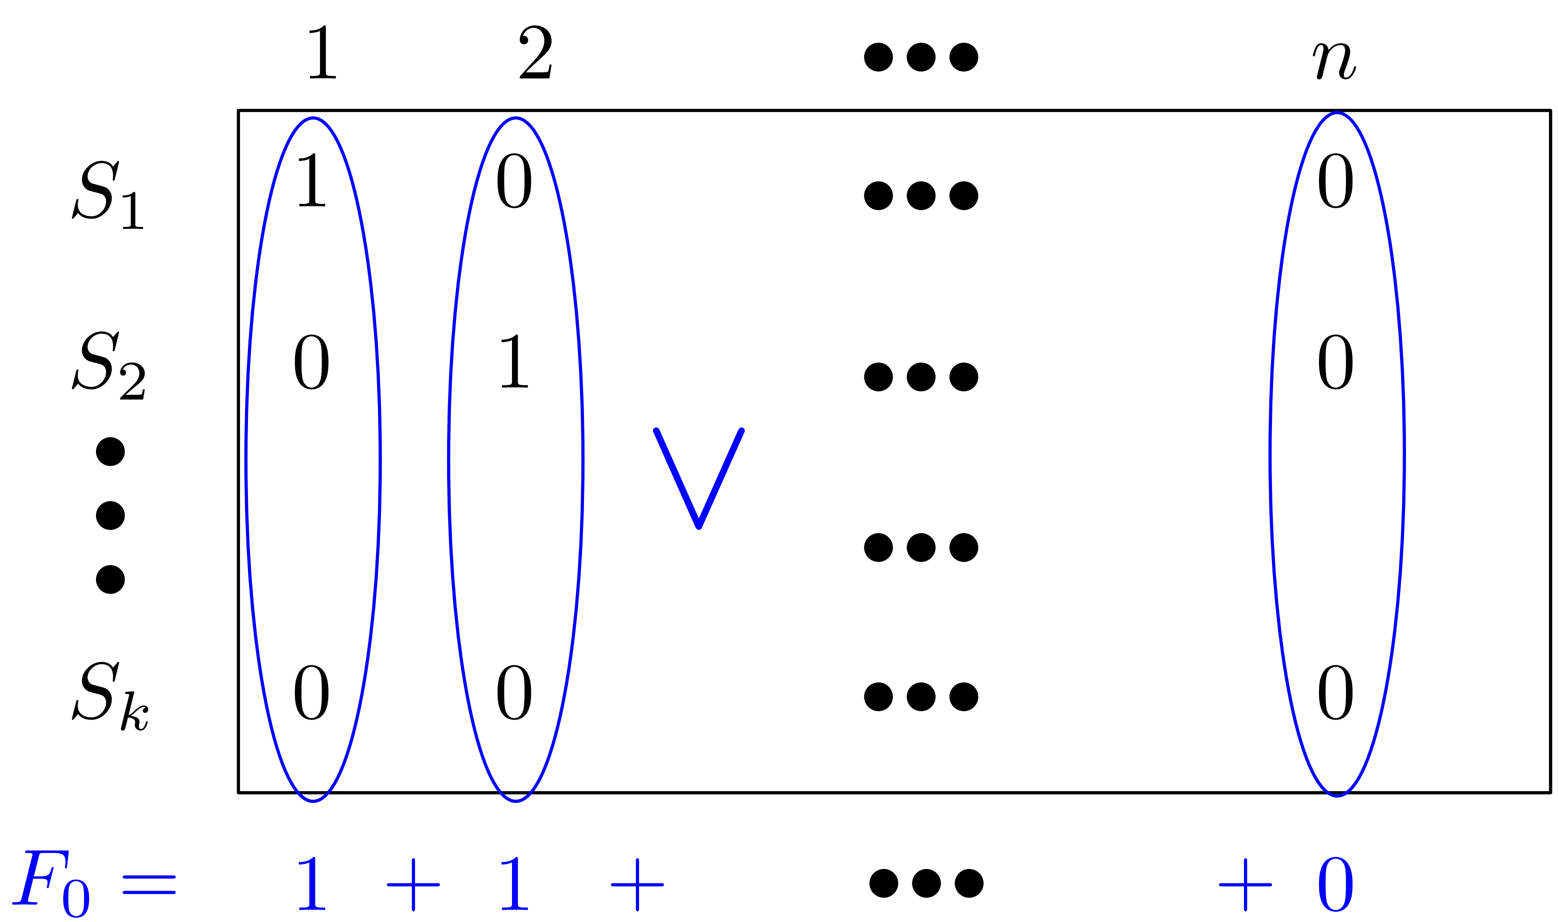

F_0 hard input distribution

Random partition

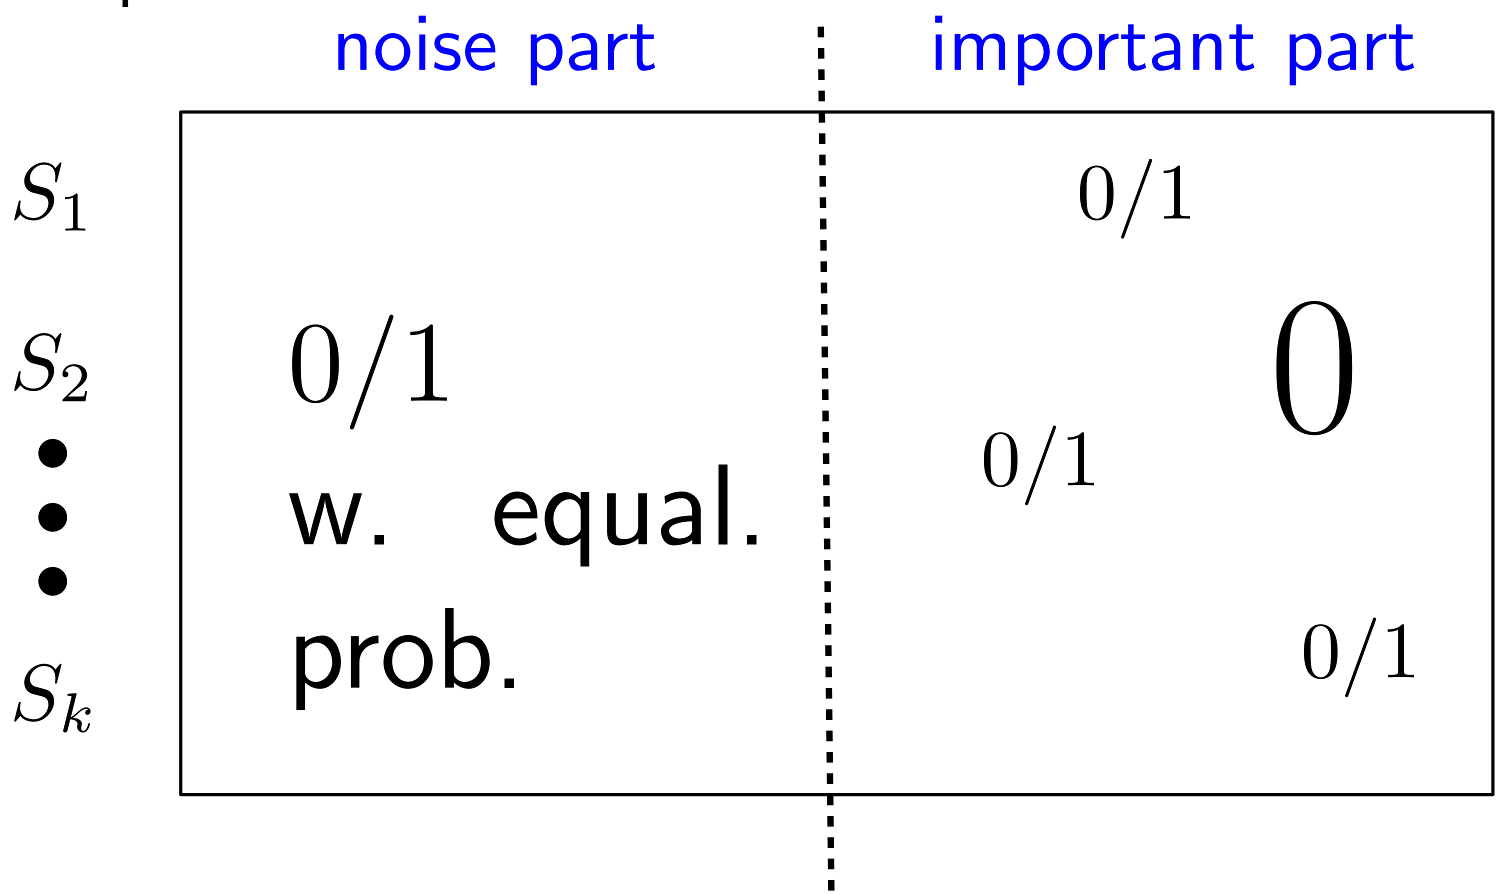

Each row in important part:

choose a random special column put "$0/1$" w. equal. prob.

rest all "0"

F_0 hard input distribution

Random partition

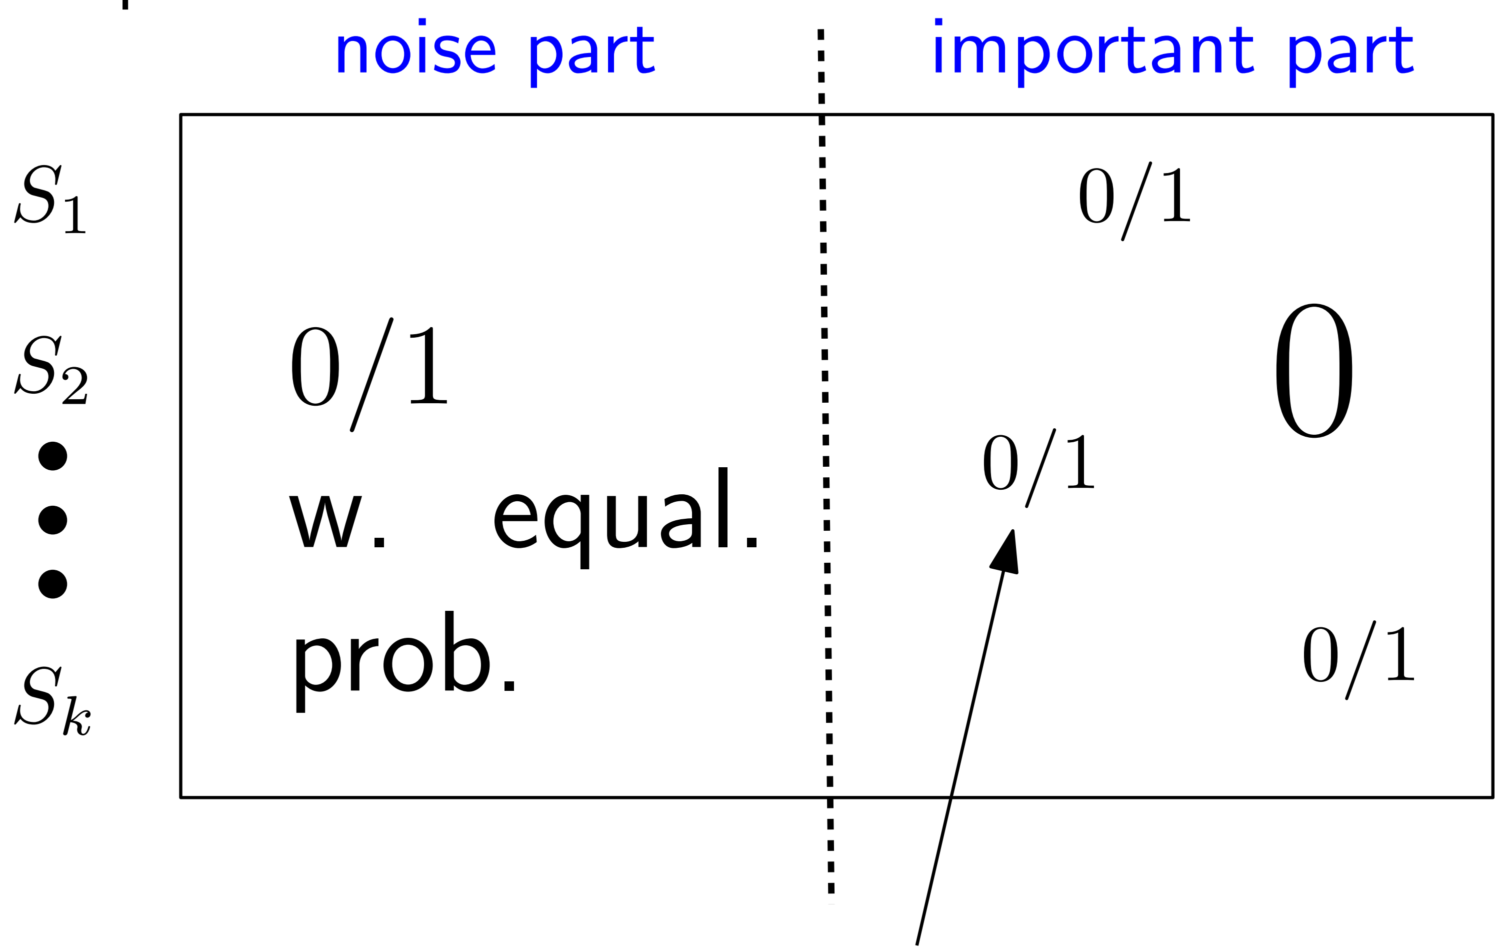

Each row in important part:

choose a random special column put "$0/1$" w. equal. prob.

rest all "0"

Let Z_i be the value in the special column in i-th row.

$F_0 \approx$ (# columns of noise part) $+ \sum_i Z_i$

The proof framework

Step 1: Find two modular problems k-GAP-MAJ and 2-DISJ of simpler structures s.t.

F_0 can be thought as a composition of them.

(Different from traditional direct-sum)

Step 2: Analyze the complexities of k-GAP-MAJ and 2-DISJ.

Step 3: Compose two modular problems so that:

$$\text{Complexity}(F_0) = \text{Complexity}(k\text{-GAP-MAJ}) \times \text{Complexity}(2\text{-DISJ}).$$

k-GAP-MAJ

k sites each holds a bit Z_i chosen uniform at random from $\{0,1\}$

Goal: compute

$$k\text{-GAP-MAJ}(Z_1, Z_2, \ldots, Z_k) = \begin{cases} 0, & \text{if } \sum_{i \in [k]} Z_i \leq k/2 - \sqrt{k}, \\ 1, & \text{if } \sum_{i \in [k]} Z_i \geq k/2 + \sqrt{k}, \\ \text{don't care}, & \text{otherwise,} \end{cases}$$

k-GAP-MAJ

k sites each holds a bit Z_i chosen uniform at random from $\{0,1\}$

Goal: compute

$$k\text{-GAP-MAJ}(Z_1, Z_2, \ldots, Z_k) = \begin{cases} 0, & \text{if } \sum_{i \in [k]} Z_i \leq k/2 - \sqrt{k}, \\ 1, & \text{if } \sum_{i \in [k]} Z_i \geq k/2 + \sqrt{k}, \\ \text{don't care}, & \text{otherwise}, \end{cases}$$

Lemma: Any protocol Π that computes k-GAP-MAJ correctly w.p. 0.9999 has to learn $\Omega(k)$ Z_i's well, that is,

$$H(Z_i \mid \Pi) \leq H_b(1/100).$$

Set disjointness (2-DISJ)

$x \cap y = \emptyset?$

$x \subseteq \{0, 1, \ldots, n\}$

$y \subseteq \{0, 1, \ldots, n\}$

Set disjointness (2-DISJ)

$x \cap y = \emptyset$?

$x \subseteq \{0, 1, \ldots, n\}$ $y \subseteq \{0, 1, \ldots, n\}$

Exists a hard distribution μ, under which

$|X \cap Y| = 1$ (YES instance) w.p. $1/2$ and

$|X \cap Y| = 0$ (NO instance) w.p. $1/2$.

Lemma: Any protocol Π that computes 2-DISJ correctly w.p. 0.99 under distribution μ communicates at least $\Omega(n)$ bits.

[Razborov '90, BJKS '04]

Next step: Compose k-GAP-MAJ with 2-DISJ

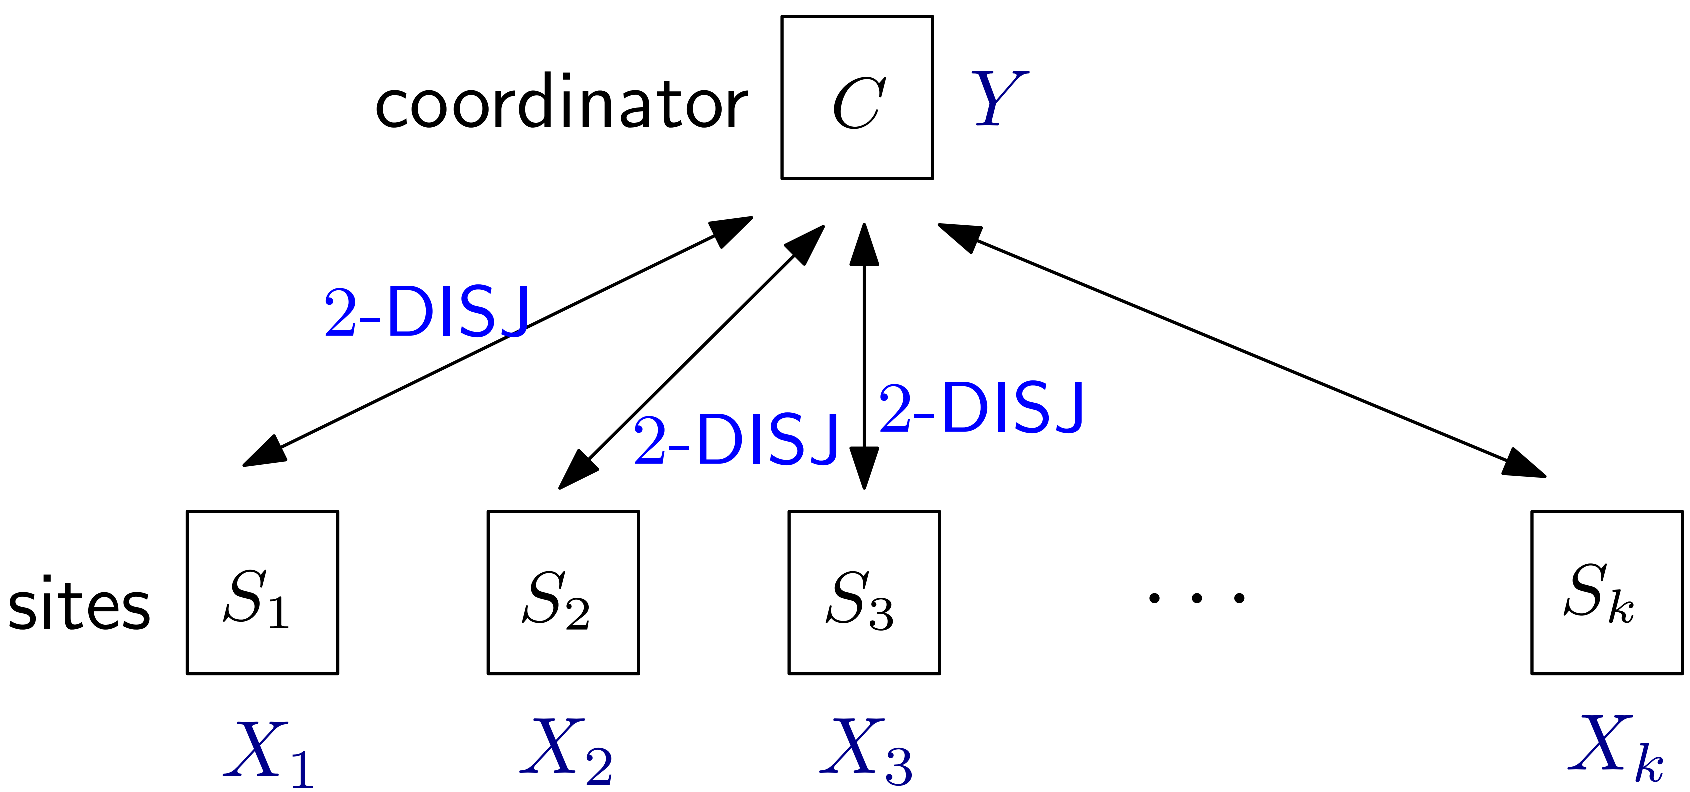

$(X_i, Y) \sim \mu$ for each $i = 1, \ldots, k$

Next step: Compose k-GAP-MAJ with 2-DISJ

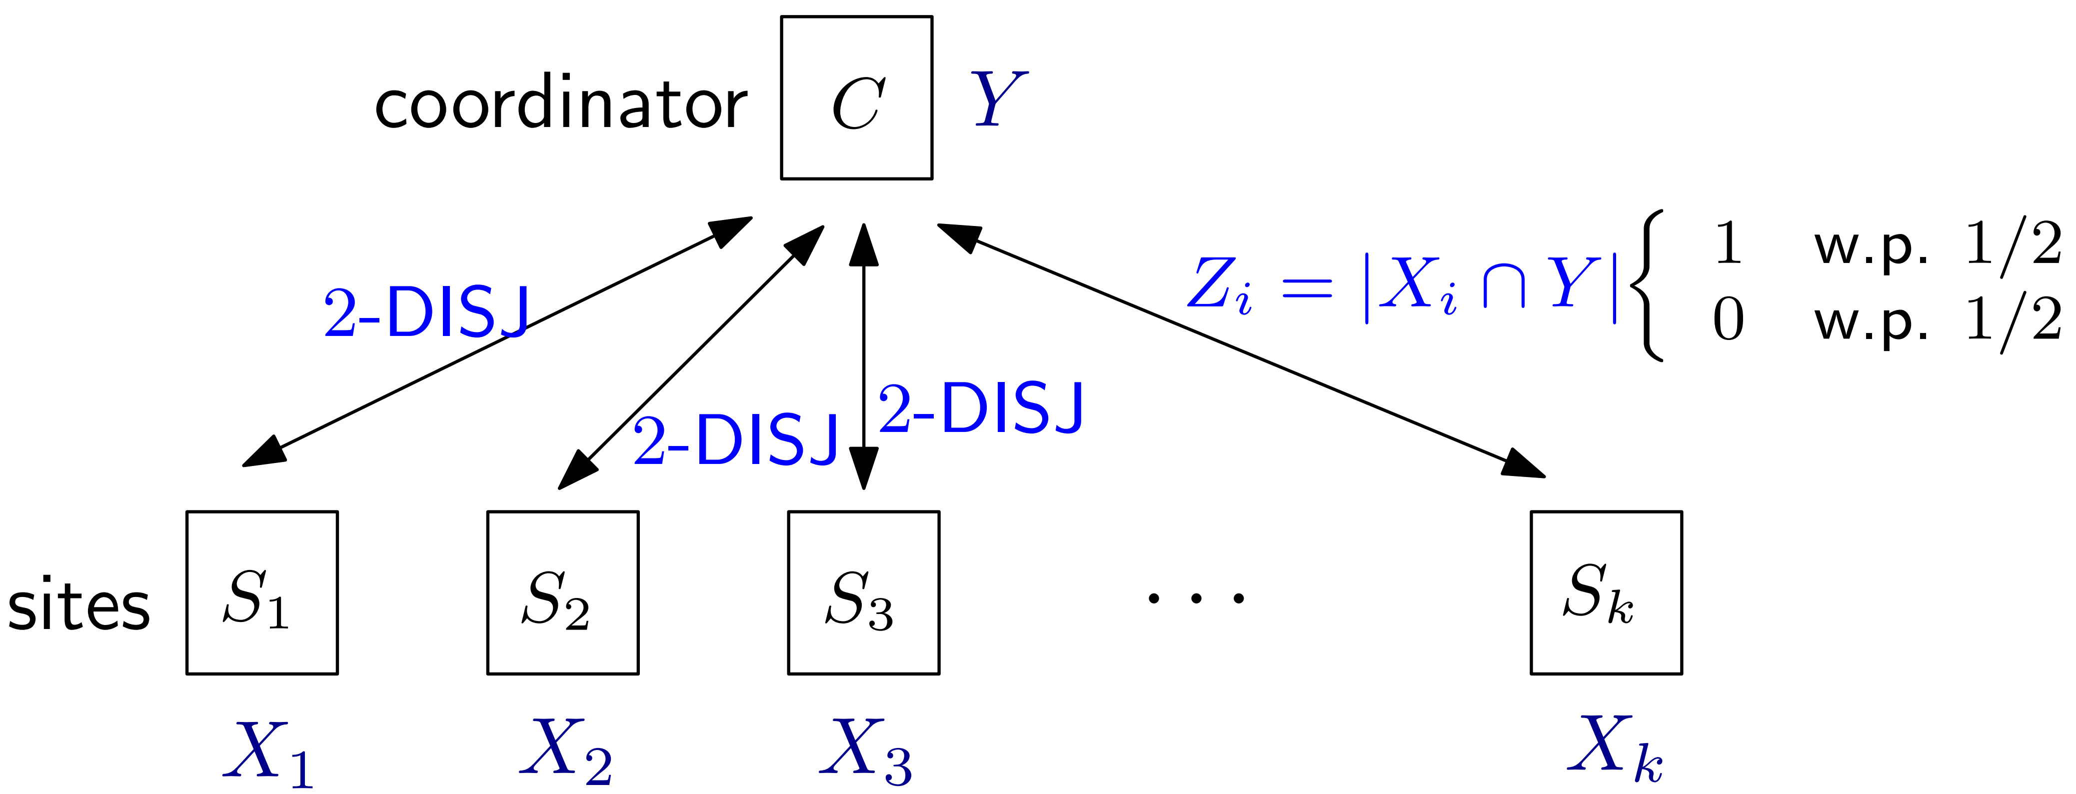

$(X_i, Y) \sim \mu$ for each $i = 1, \ldots, k$

Next step: Compose k-GAP-MAJ with 2-DISJ

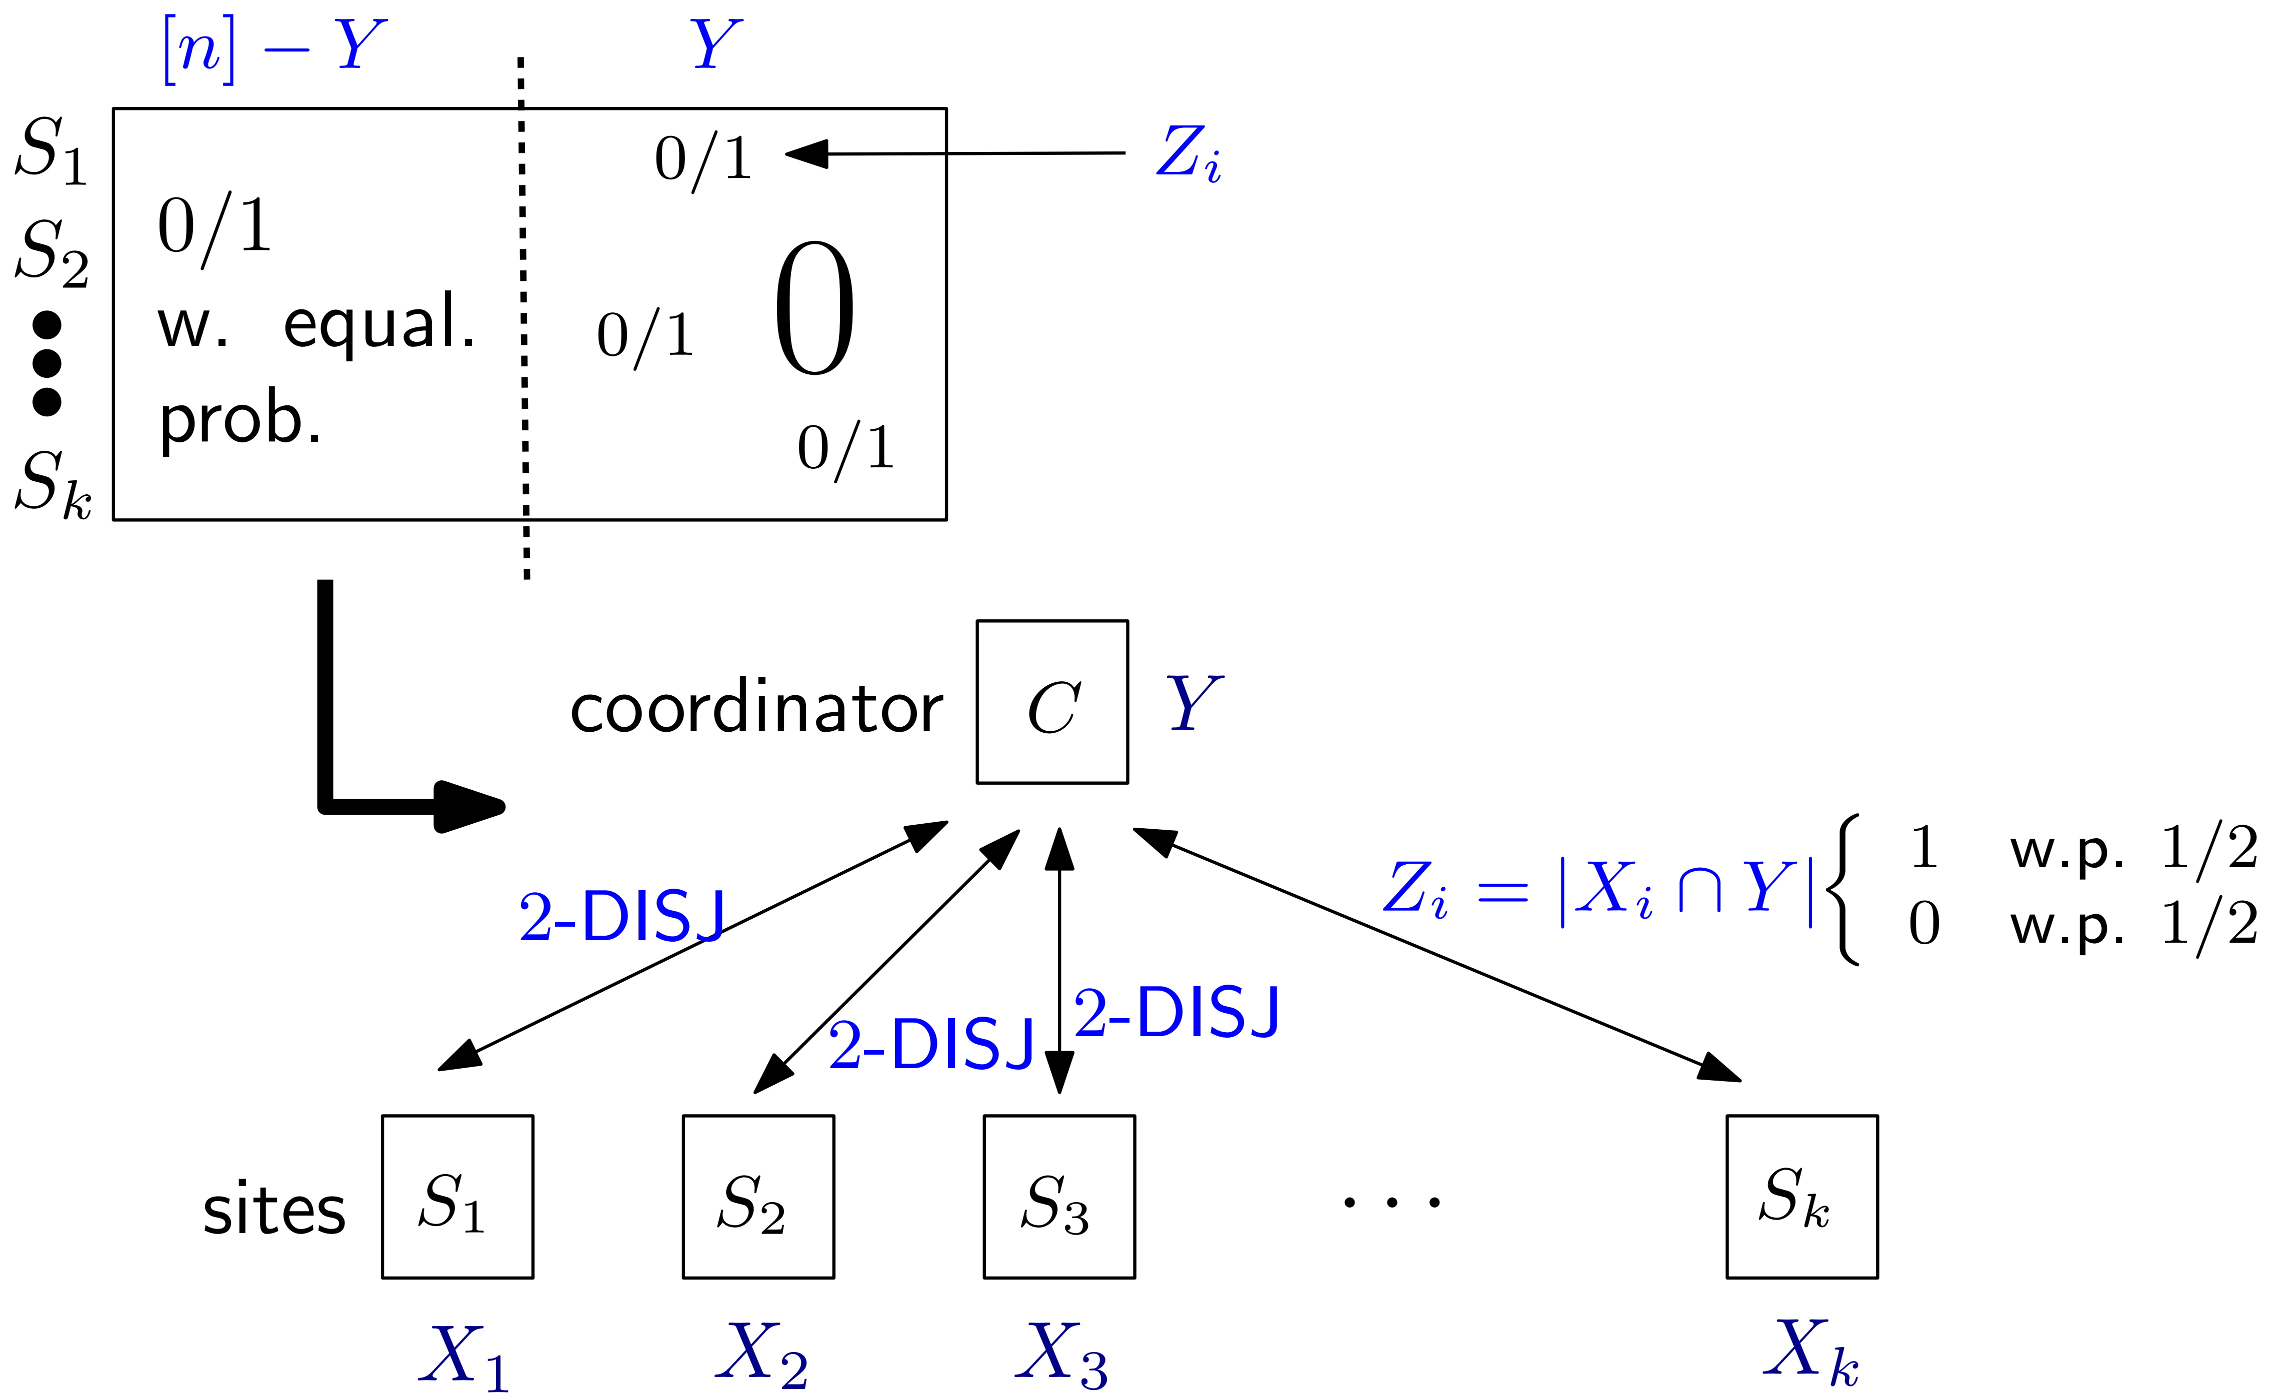

$(X_i, Y) \sim \mu$ for each $i = 1, \ldots, k$

Next step: Compose k-GAP-MAJ with 2-DISJ

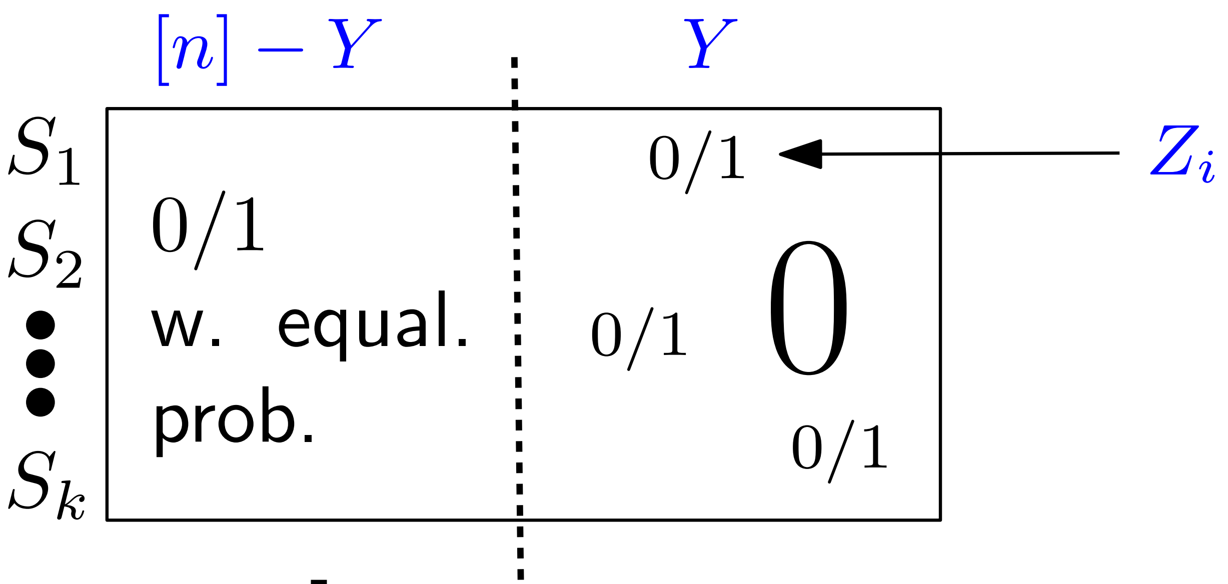

$$F_0(X_1, X_2, \ldots, X_k) \approx |[n]-Y| + \sum_i Z_i$$

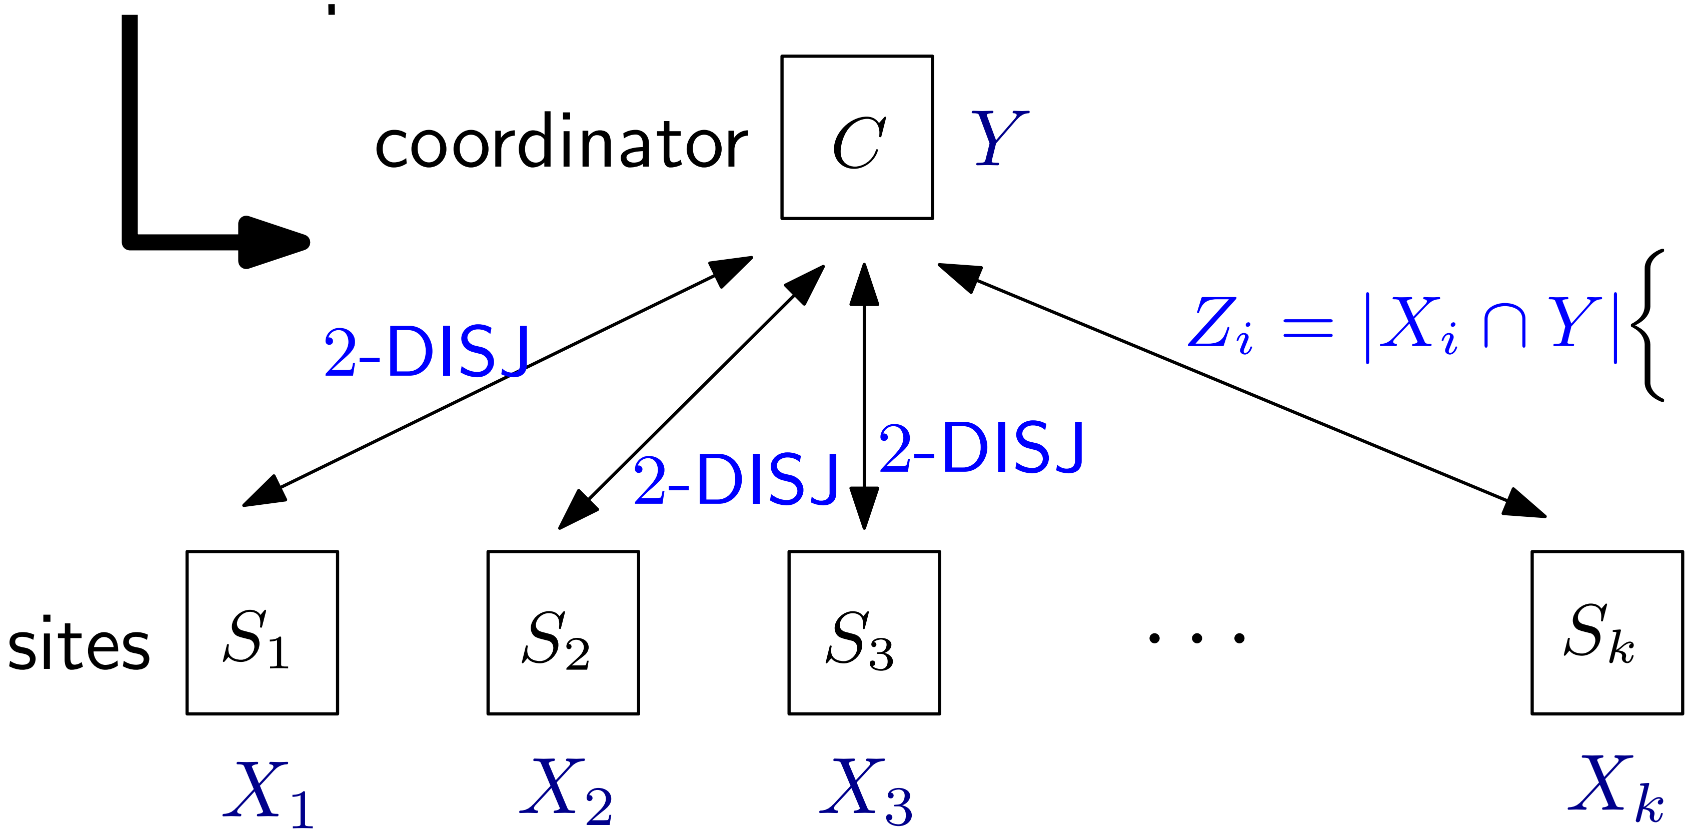

$$Z_i = |X_i \cap Y| \begin{cases} 1 & \text{w.p. } 1/2 \\ 0 & \text{w.p. } 1/2 \end{cases}$$

$(X_i, Y) \sim \mu$ for each $i = 1, \ldots, k$

Next step: Compose k-GAP-MAJ with 2-DISJ

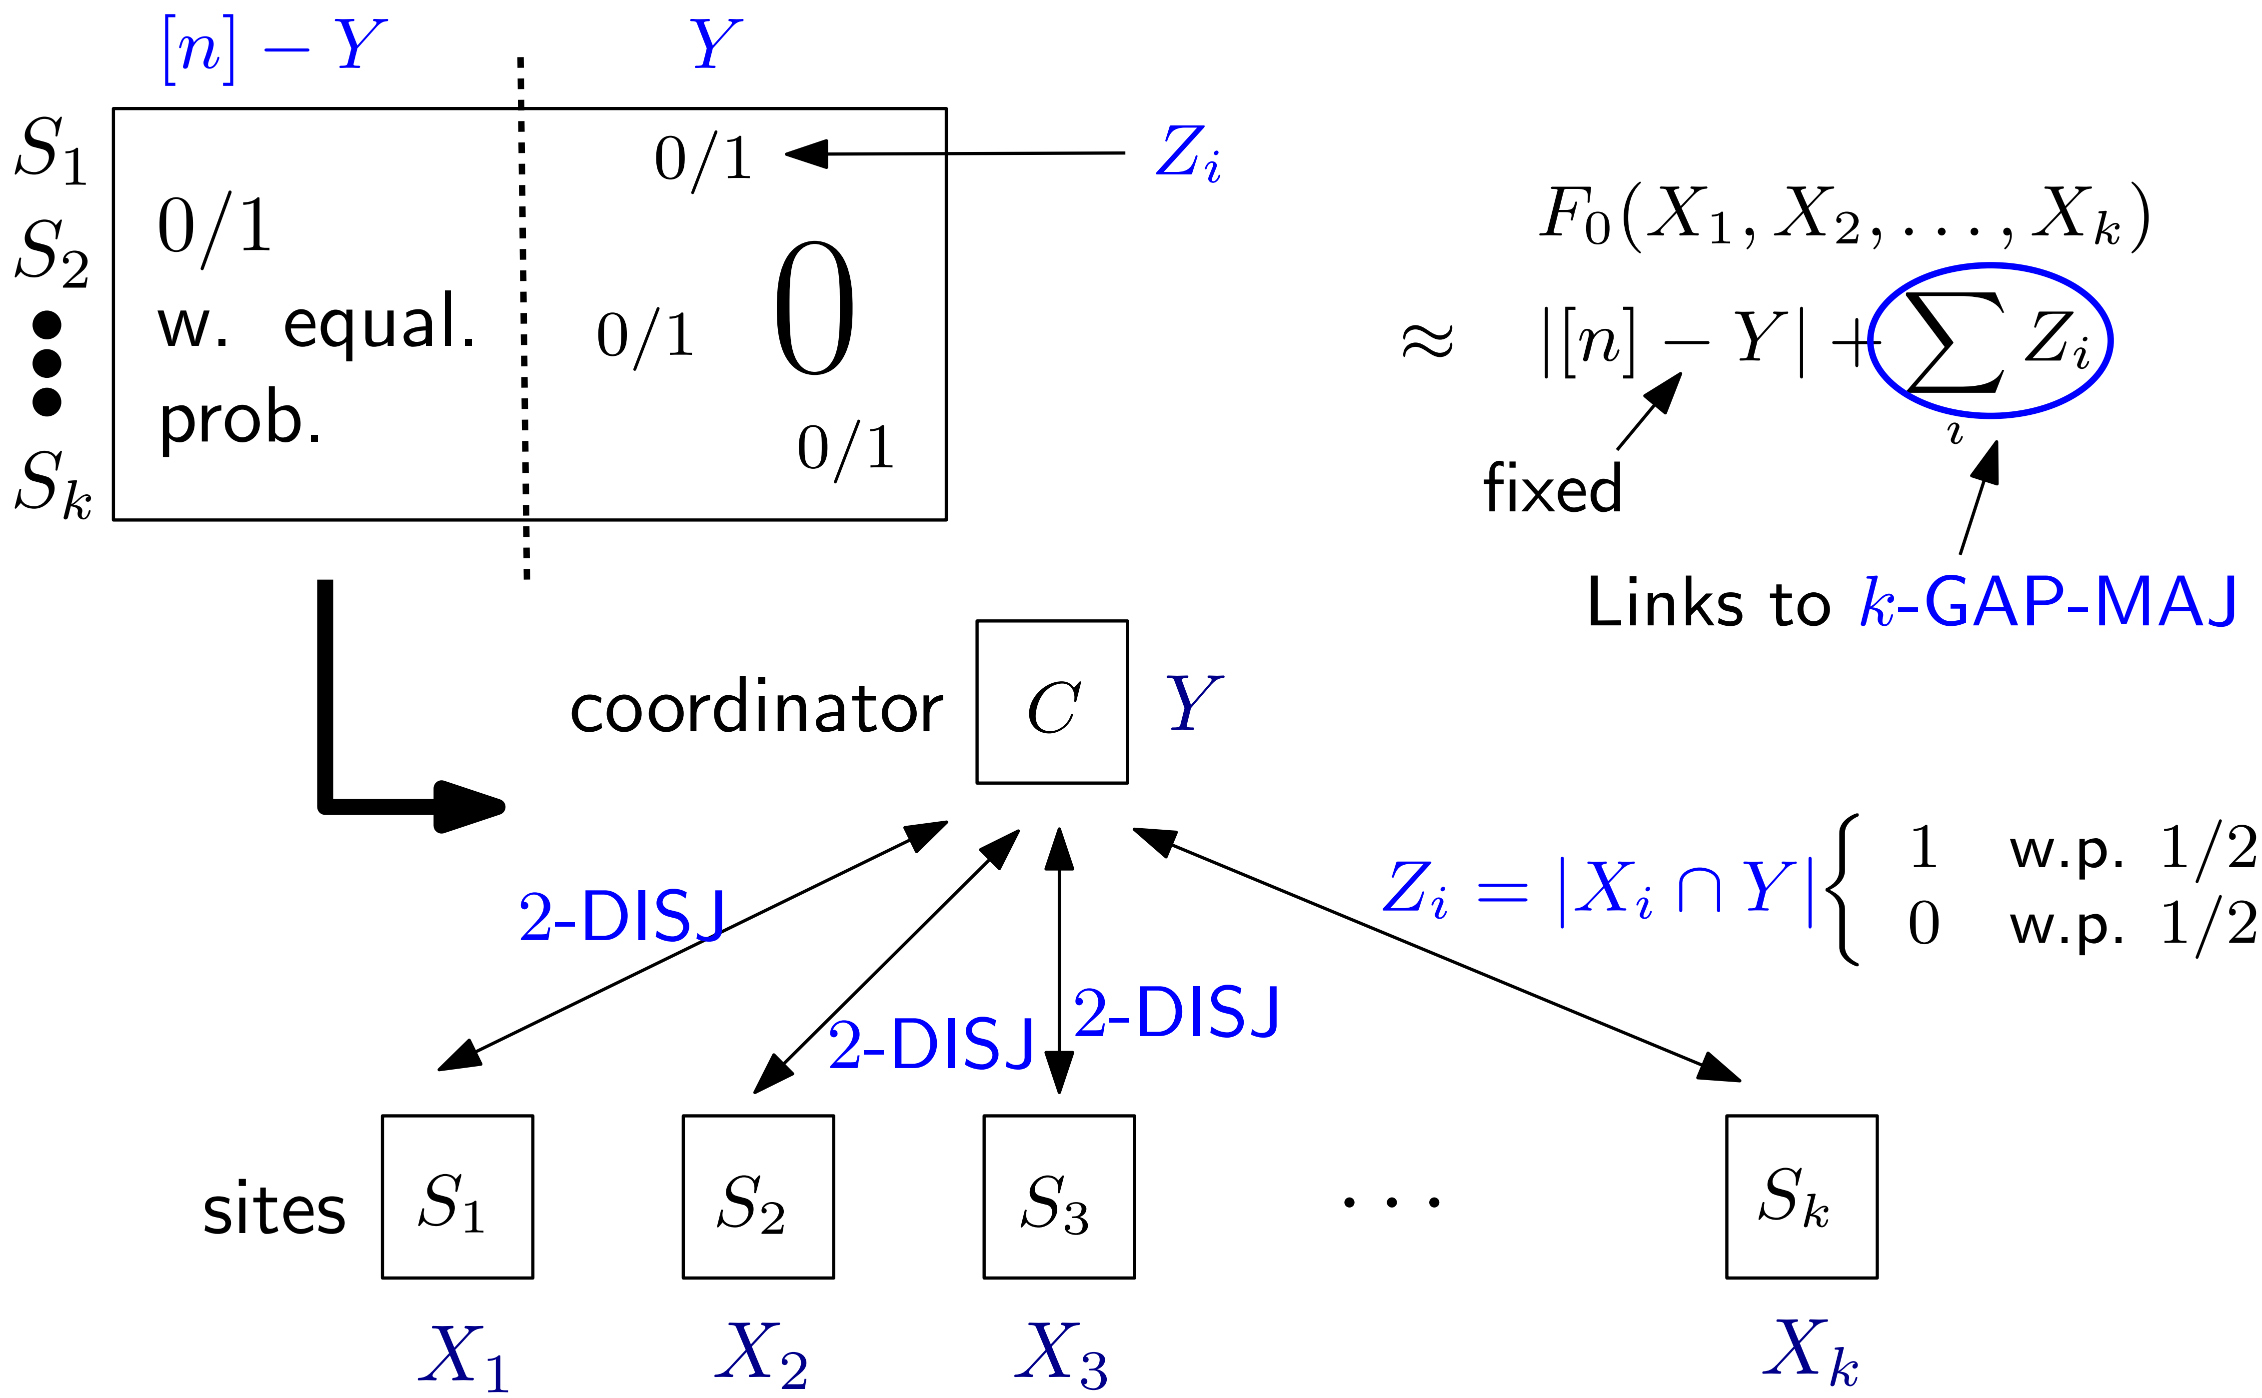

$(X_i, Y) \sim \mu$ for each $i = 1, \ldots, k$

The proof

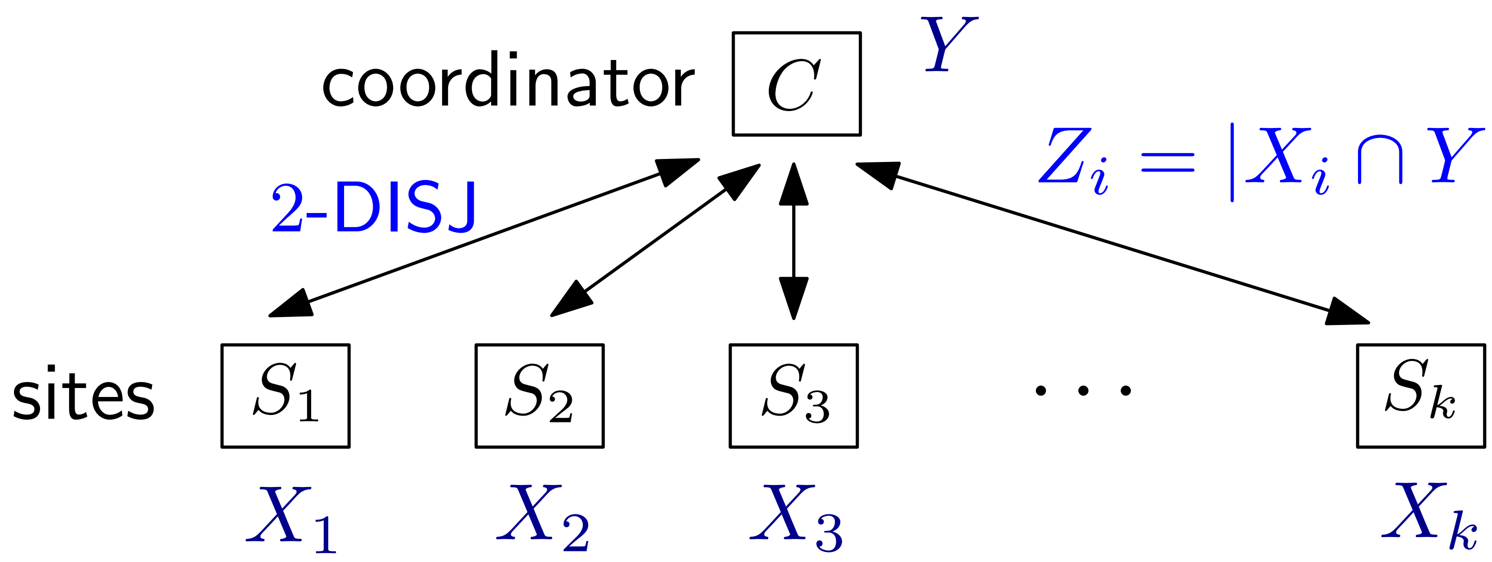

$$Z_i = |X_i \cap Y| \begin{cases} 1 & \text{w.p. } 1/2 \\ 0 & \text{w.p. } 1/2 \end{cases}$$

$F_0(X_1, X_2, \ldots, X_k) \iff k$-GAP-MAJ$(Z_1, Z_2, \ldots, Z_k)$

$(Z_i = |X_i \cap Y|)$

$\Longrightarrow$ learn $\Omega(k)$ Z_i's well

$\Longrightarrow$ need $\Omega(nk)$ bits

(observe $Z_i | \Pi$ are independent, learning each Z_i well needs $\Omega(n)$ bits, by 2-DISJ)

The proof

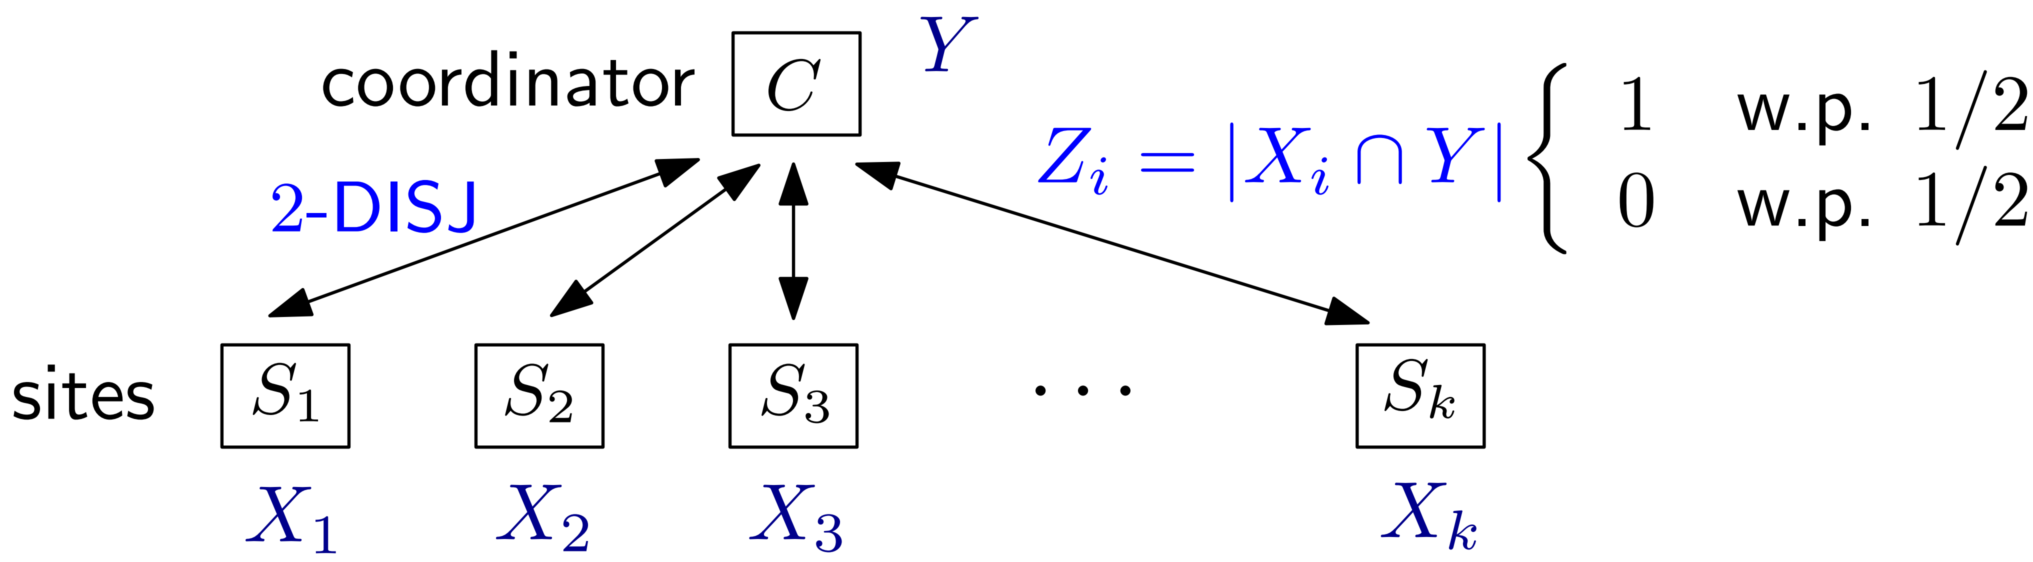

$F_0(X_1, X_2, \ldots, X_k) \iff k$-GAP-MAJ$(Z_1, Z_2, \ldots, Z_k)$

$(Z_i = |X_i \cap Y|)$

$\implies$ learn $\Omega(k)$ Z_i's well

$\implies$ need $\Omega(nk)$ bits

(observe $Z_i|\Pi$ are independent, learning each Z_i well needs $\Omega(n)$ bits, by 2-DISJ)

For the reduction set $n = \Theta(1/\varepsilon^2)$ we get $\Omega(k/\varepsilon^2)$ LB.

Q.E.D.

Our new technique: composition

Step 1: Find two (or more) modular problems A, B of simpler structures s.t.

the original problem can be thought as a composition of them.

Step 2: Analyze the complexities of A, B.

Step 3: Compose modular problems A, B so that:

Complexity(original problem) $=$

Complexity$(A)\times$ Complexity(B).

The F_2 problem

k sites each holds a set X_i $(i \in \{1, 2, \ldots, k\})$.

Goal: compute $F_2(\cup_{i=1}^{k} X_i)$ up to a $(1+\varepsilon)$-approximation.

Previous UB: $\tilde{O}(k^2/\varepsilon + k^{1.5}/\varepsilon^3)$

Our UB: $\tilde{O}(k/\text{poly}(\varepsilon))$, one way protocol

Previous LB: $\Omega(k)$

Our LB: $\tilde{\Omega}(k/\varepsilon^2)$. Holds in blackboard model. Tight in static case.

The F_2 problem

k sites each holds a set X_i $(i \in \{1, 2, \ldots, k\})$.

Goal: compute $F_2(\cup_{i=1}^{k} X_i)$ up to a $(1 + \varepsilon)$-approximation.

Previous UB: $\tilde{O}(k^2/\varepsilon + k^{1.5}/\varepsilon^3)$

Our UB: $\tilde{O}(k/\text{poly}(\varepsilon))$, one way protocol

Previous LB: $\Omega(k)$

Our LB: $\tilde{\Omega}(k/\varepsilon^2)$. Holds in blackboard model. Tight in static case.

LB proof ideas overview:

Same framework, choose two modular problems

- Gap-Hamming
- k-DISJ

The F_2 problem

k sites each holds a set X_i $(i \in \{1, 2, \ldots, k\})$.

Goal: compute $F_2(\cup_{i=1}^k X_i)$ up to a $(1+\varepsilon)$-approximation.

Previous UB: $\tilde{O}(k^2/\varepsilon + k^{1.5}/\varepsilon^3)$

Our UB: $\tilde{O}(k/\text{poly}(\varepsilon))$, one way protocol

Previous LB: $\Omega(k)$

Our LB: $\tilde{\Omega}(k/\varepsilon^2)$. Holds in blackboard model. Tight in static case.

LB proof ideas overview:
Same framework, choose two modular problems

- Gap-Hamming
- k-DISJ

– Compose in a different way to prove a LB for F_2

– Heavy use of information cost

Two modular problems

- 2-**party Gap-Hamming**: Alice has $X = \{X_1, X_2, \ldots, X_{1/\varepsilon^2}\}$, Bob has $Y = \{Y_1, Y_2, \ldots, Y_{1/\varepsilon^2}\}$. They want to compute:

$$\textbf{GHD}(X,Y) = \begin{cases} 0, & \text{if } \sum_{i \in [1/\varepsilon^2]} X_i \oplus Y_i \le 1/2\varepsilon^2 - 1/\varepsilon, \\ 1, & \text{if } \sum_{i \in [1/\varepsilon^2]} X_i \oplus Y_i \ge 1/2\varepsilon^2 + 1/\varepsilon, \\ *, & \text{otherwise,} \end{cases}$$

  Solving it w.r.t. uniform distribution needs to reveal $\Omega(1/\varepsilon^2)$ bits of the input

Two modular problems

- 2-**party Gap-Hamming**: Alice has $X = \{X_1, X_2, \ldots, X_{1/\varepsilon^2}\}$, Bob has $Y = \{Y_1, Y_2, \ldots, Y_{1/\varepsilon^2}\}$. They want to compute:

$$\textbf{GHD}(X,Y) = \begin{cases} 0, & \text{if } \sum_{i\in[1/\varepsilon^2]} X_i \oplus Y_i \le 1/2\varepsilon^2 - 1/\varepsilon, \\ 1, & \text{if } \sum_{i\in[1/\varepsilon^2]} X_i \oplus Y_i \ge 1/2\varepsilon^2 + 1/\varepsilon, \\ *, & \text{otherwise,} \end{cases}$$

  Solving it w.r.t. uniform distribution needs to reveal $\Omega(1/\varepsilon^2)$ bits of the input

- k-**DISJ**: We have k sites $S_1, \ldots, S_k$. S_i holds a set Z_i $(|Z_i| = k^2)$. We promise that

  - either Z_i are all disjoint,
  - or they intersect on one element and the rest are all disjoint (sun-flower).

  The goal is to find out which is the case.

  Solving it w.r.t. certain distribution needs to reveal $\Omega(k)$ bits of the input

Next step: Compose Gap-Hamming with k-DISJ

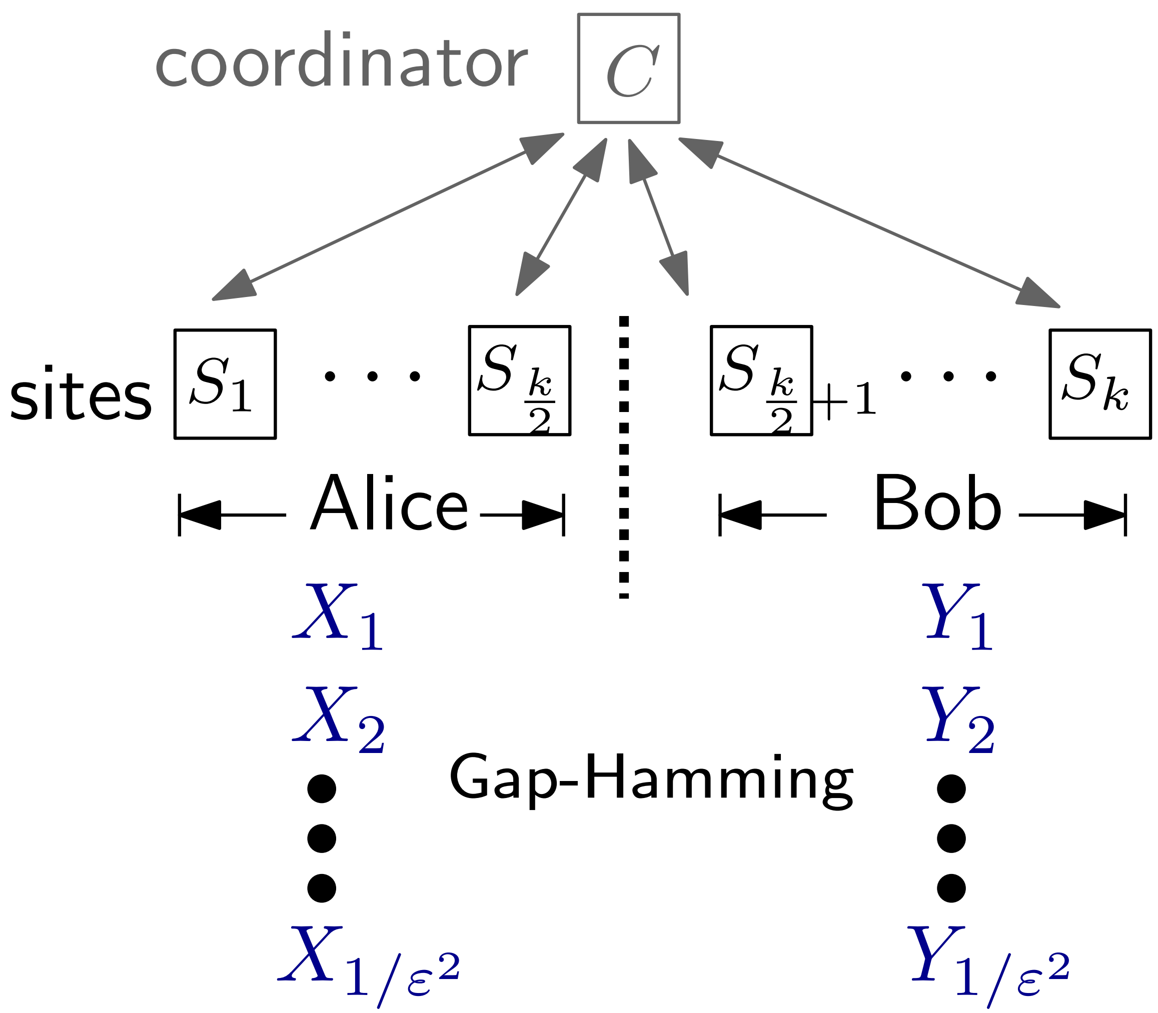

We create $2/\varepsilon^2$ $(k/2)$-DISJ instances, one for each input bit of Gap-Hamming.

Next step: Compose Gap-Hamming with k-DISJ

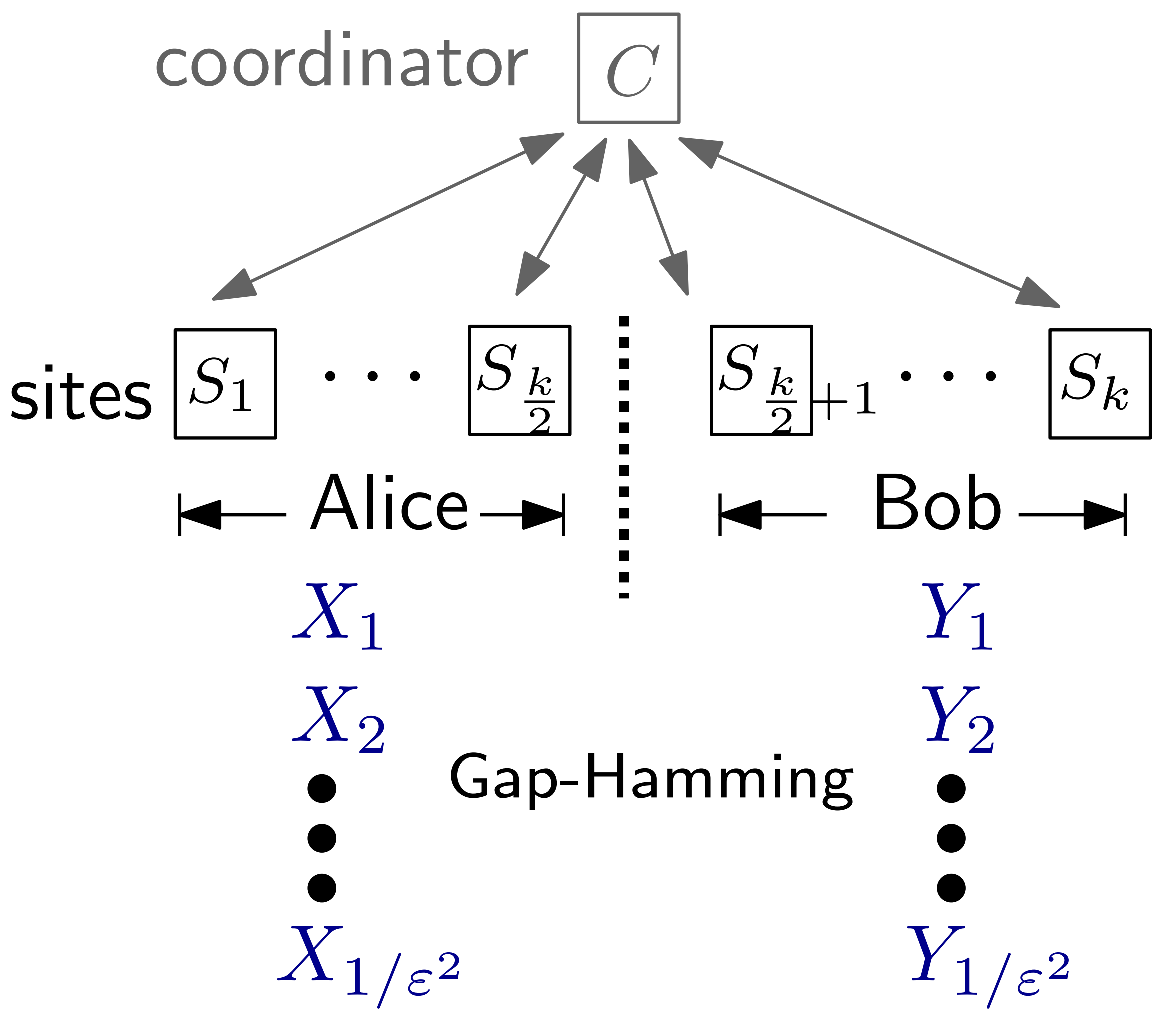

The proof

Solve $(1+\varepsilon)$-approx F_2

$\Leftrightarrow$ Solve Gap-Hamming (GHD)

$\Leftrightarrow$ Learn $\Omega(1/\varepsilon^2)$ input bits of GHD

- Learning each bit of GHD needs to solve an instance of $(k/2)$-DISJ
- Solving each $(k/2)$-DISJ has to reveal $\Omega(k)$ bits of the inputs

$\Rightarrow$ Reveal $\Omega(k/\varepsilon^2)$ bits in total

We create $2/\varepsilon^2$ $(k/2)$-DISJ instances,
one for each input bit of Gap-Hamming.

Q.E.D.

F_p $(p > 1)$ upper bounds, a quick glance

- Previous UB: $\tilde{O}(k^{2p+1} n^{1-2/p} / \text{poly}(\varepsilon))$

  Our UB: $\tilde{O}(k^{p-1} / \text{poly}(\varepsilon))$

F_p $(p > 1)$ upper bounds, a quick glance

- Previous UB: $\tilde{O}(k^{2p+1} n^{1-2/p} / \mathrm{poly}(\varepsilon))$

  Our UB: $\tilde{O}(k^{p-1} / \mathrm{poly}(\varepsilon))$

- Inspired by work on sub-sampling [Indyk-Woodruff 2005]
  New features in our protocol:

  - No AMS sketches
  - One-way protocol
  - Threshold-based sampling used to communication-efficiently implement distributed k-party heavy hitters.

F_p $(p > 1)$ upper bounds, a quick glance

- Previous UB: $\tilde{O}(k^{2p+1} n^{1-2/p} / \mathrm{poly}(\varepsilon))$

  Our UB: $\tilde{O}(k^{p-1} / \mathrm{poly}(\varepsilon))$

- Inspired by work on sub-sampling [Indyk-Woodruff 2005]
  New features in our protocol:

  - No AMS sketches
  - One-way protocol
  - Threshold-based sampling used to communication-efficiently implement distributed k-party heavy hitters.

- We suspect it can have more applications, as IW05 did for streaming model. e.g., for distributed EMD, distributed l_p-sampling, etc.

Conclusion and future work

- We obtained: (almost) tight bounds for frequency moments, heavy hitters, quantile and entropy in the distributed streaming model.

Conclusion and future work

- We obtained: (almost) tight bounds for frequency moments, heavy hitters, quantile and entropy in the distributed streaming model.
- Future work
  - F_2 is not tight in terms of ε

Conclusion and future work

- We obtained: (almost) tight bounds for frequency moments, heavy hitters, quantile and entropy in the distributed streaming model.

- Future work

  - F_2 is not tight in terms of ε

  - Generalize the model: consider the network topology; items go into multiple sites, ...

Conclusion and future work

- We obtained: (almost) tight bounds for frequency moments, heavy hitters, quantile and entropy in the distributed streaming model.
- Future work
  - F_2 is not tight in terms of ε
  - Generalize the model: consider the network topology; items go into multiple sites, ...
  - Beyond statistical problems
    - Geometry problems: range-counting, extent measures, . . .
    - Graph problems

The end

THANK YOU

Questions?