When Distributed Computation is Communication Expensive

David P. Woodruff
IBM Research Almaden

Qin Zhang
Indiana University Bloomington

DISC'13,
October, 2013

The coordinator model

- **The coordinator model**: We have k machines (sites) and one central server.

  – Each site has a 2-way communication channel with the server.
  – Each site has a piece of data $x_i (i \in [k])$.
  – Goal is to compute $f(x_1, \ldots, x_i)$ together via communication, for some function f.

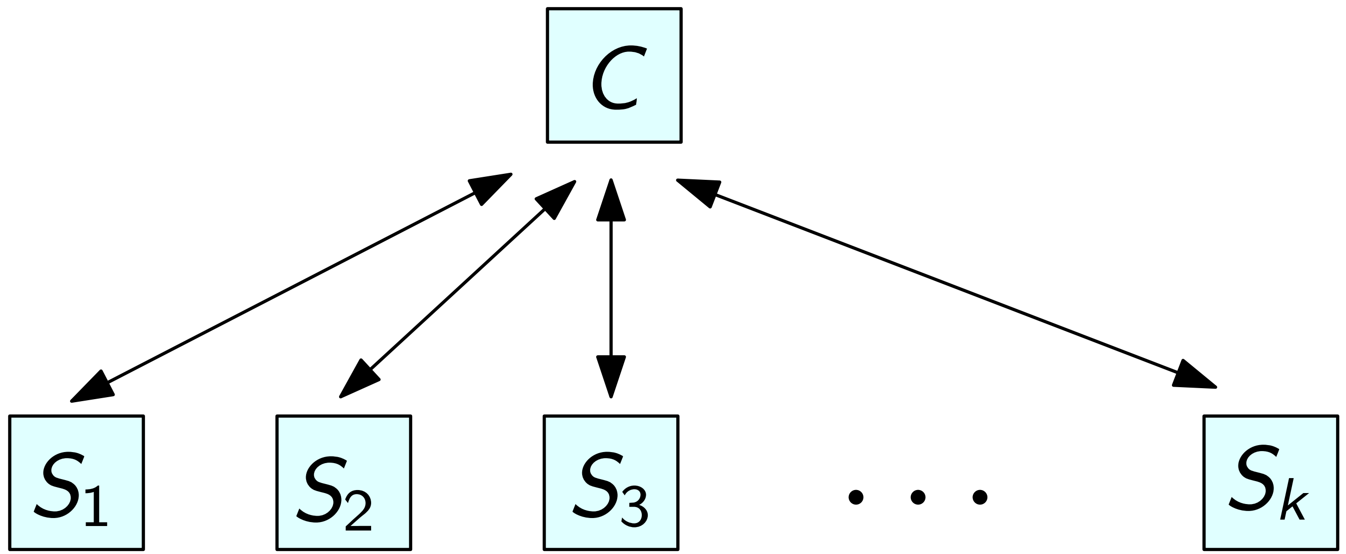

Motivation

Communication → energy, bandwidth, . . .

Sensor networks

Data in the cloud

Motivation

Communication → energy, bandwidth, . . .

Sensor networks

Data in the cloud

Abstraction

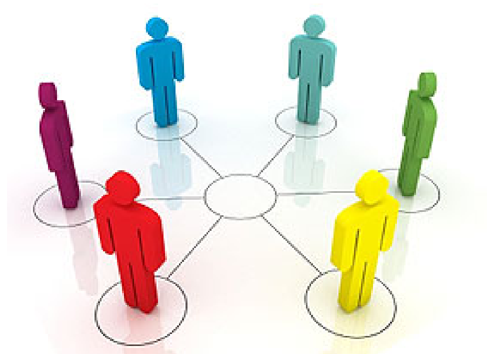

Motivation

Communication → energy, bandwidth, . . .

Data in the cloud

Sensor networks

Abstraction

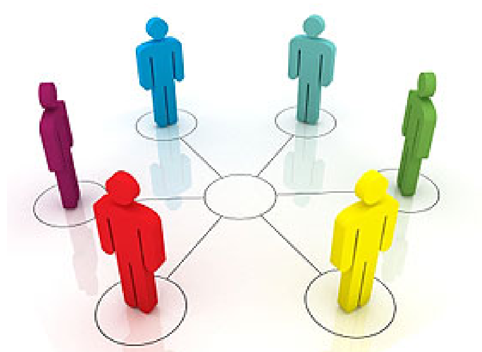

=

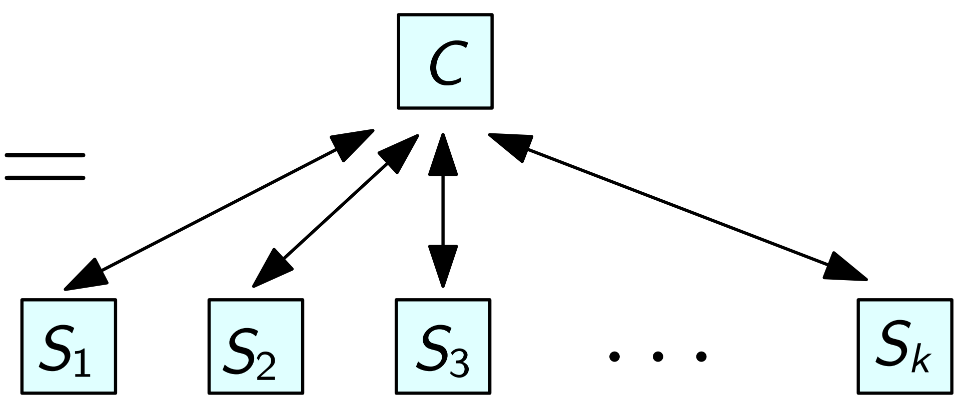

Conceptual messages

If we want communication efficient protocols, then for many problems/functions f,

- Allow an approximate solution.
  – For a problem whose output is a single numerical value x, approximation means outputting $\tilde{x} \in [(1-\epsilon)x, (1+\epsilon)x]$.
  – For a problem whose output is YES/NO, approximation means outputting YES if the input is close to have the property, and NO otherwise.

- Use well-designed input layouts.

- Explore prior distributional properties of the input dataset.

Our results

	With duplication		Without duplication	
Problem	LB	UB	LB	UB
F_0	$\tilde{\Omega}(kF_0)$	$\tilde{O}(k(F_0 + \log n))$	–	–
ℓ_∞	$\tilde{\Omega}(\min\{k, \ell_\infty\}n)$	$\tilde{O}(\min\{k, \ell_\infty\}n)$	–	–
degree	$\tilde{\Omega}(kd_v)$	$O(kd_v \log n)$	$\tilde{\Omega}(k)$	$O(k \log n)$
cycle-freeness	$\tilde{\Omega}(kn)$	$\tilde{O}(kn)$	$\Omega(s)$	$\tilde{O}(s)$
connectivity	$\tilde{\Omega}(kn)$	$\tilde{O}(kn)$	$\tilde{\Omega}(kr)$	$\tilde{O}(kr)$
#CC	$\tilde{\Omega}(kn)$	$\tilde{O}(kn)$	$\tilde{\Omega}(kr)$	$\tilde{O}(kr)$
bipartiteness	$\tilde{\Omega}(kn)$	$\tilde{O}(kn)$	$\tilde{\Omega}(kr)$	$\tilde{O}(kr)$
triangle-freeness	$\tilde{\Omega}(km)$	$\tilde{O}(km)$	$\Omega(m)$	$\tilde{O}(m)$

– F_0: # distinct elements.
– ℓ_∞: the element with the maximum frequency.
– #CC: # connected components

All results are in terms of number of bits. Our lower bounds hold for randomized protocols with successful probability $2/3$.

For F_0 and ℓ_∞, n denotes the size of the universe. For graph problems, n denotes # vertices and m denotes # edges. d_v is the degree of the vertex v. Let $r = \min\{n, m/k\}$, $s = \min\{n, m\}$. $\tilde{O}, \tilde{\Omega}$ hides log factors.

What do these results imply?

Case study 1: F_0 – the value of approximation

- F_0: Compute # distinct elements of the union of k data sets

Case study 1: F_0 – the value of approximation

- F_0: Compute $\#$ distinct elements of the union of k data sets
- Our lower bound for exact computation is $\Omega(kF_0/\log k)$ bits.
- We can implement a streaming algorithm for $(1+\epsilon)$-approximation F_0 in [Kane, Nelson, Woodruff. 2010] in the coordinator model using $O(k(\log n + 1/\epsilon^2 \log 1/\epsilon))$ bits of communication.

Case study 1: F_0 – the value of approximation

- F_0: Compute # distinct elements of the union of k data sets
- Our lower bound for exact computation is $\Omega(kF_0/\log k)$ bits.
- We can implement a streaming algorithm for $(1+\epsilon)$-approximation F_0 in [Kane, Nelson, Woodruff. 2010] in the coordinator model using $O(k(\log n + 1/\epsilon^2 \log 1/\epsilon))$ bits of communication.
- In a typical setting, we could have $\epsilon = 0.01$, $n = 10^9$ and $k = 1000$. Then $O(k(\log n + 1/\epsilon^2 \log 1/\epsilon)) = 6.6 \times 10^7$ bits, while $\Omega(kF_0/\log k) = 10^9$ bits even when $F_0 = n/100$.

Case study 2: Connectivity – the value of input layout

- Connectivity: Test if a graph is connected.

Case study 2: Connectivity – the value of input layout

- Connectivity: Test if a graph is connected.

- If edges are partitioned arbitrarily in the k sites edge-partition, we get a lower bound of $\Omega(kn/\log k)$ bits.

- If each node together with all its adjacent edges are stored in one of the k sites (node-partition), then we can adapt a sketching algorithm by Ahn, Guha, McGregor (2012) using $O(k + n \operatorname{poly}\log n)$ bits.

Case study 3: Diameter – the value of approximation

- Diameter: Compute the diameter of the graph.

Case study 3: Diameter – the value of approximation

- Diameter: Compute the diameter of the graph.

- A conjecture in our paper (very recently proved by Braverman et al.): The communication cost of exact computation in the coordinator model is $\Omega(km)$ bits (allow edge duplication).

- There exists an algorithm (a distributed implementation of an algorithm Dor, Halperin, and Zwick, 2000) with cost $\tilde{O}(kn^{1.5})$ if an approximation of additive 2 is allowed.

How to prove these results?

How to prove these results?

Use reductions from a meta-problem THRESH (defined later)

Basics on communication complexity

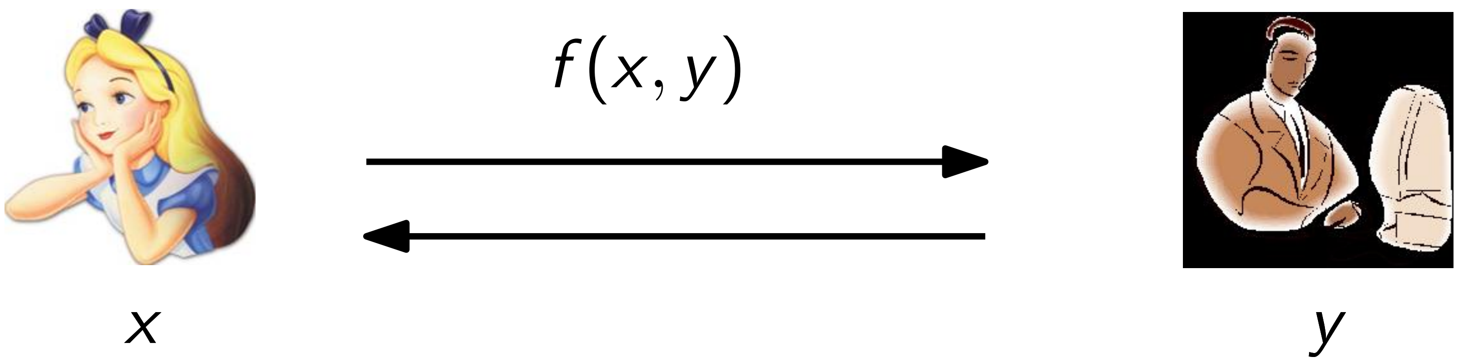

Can be easily extended to multiple players.

- R^δ: $\max_{x,y} |\Pi(x, y)|$, $|\Pi(x, y)|$ is the length of the transcript on input x, y. Π is randomized, and $\Pi(x, y) \neq f(x, y)$ w.pr. at most δ for any (x, y).
- D_μ^δ: $\max_{x,y} |\Pi(x, y)|$. Π is deterministic, and $\Pi(x, y) \neq f(x, y)$ for at most a δ fraction of (x, y) under distribution μ.
- ED_μ^δ: $\mathrm{E}_{(x,y)\sim\mu}|\Pi(x, y)|$. Π is deterministic, and $\Pi(x, y) \neq f(x, y)$ for at most a δ fraction of (x, y) under distribution μ.

Easy direction of Yao's Lemma: $R^\delta(f) \geq \max_\mu D_\mu^\delta(f)$.

A direct-sum type theorem in the coordinator model

$f : \mathcal{X} \times \mathcal{Y} \rightarrow \{0, 1\}$

μ is a distribution over $\mathcal{X} \times \mathcal{Y}$.

$f^k_{\mathsf{OR}} : \mathcal{X}^k \times \mathcal{Y} \rightarrow \{0, 1\}$ is the problem of computing $f(x_1, y) \vee f(x_2, y) \vee \ldots \vee f(x_k, y)$ in the coordinator model, where P_i has input $x_i \in \mathcal{X}$ for each $i \in [k]$, and the coordinator has $y \in \mathcal{Y}$.

ν is a distribution on $\mathcal{X}^k \times \mathcal{Y}$: First pick $(X_1, Y) \sim \mu$, and then pick $X_2, \ldots, X_k$ from the conditional distribution $\mu \mid Y$.

A direct-sum type theorem in the coordinator model

$f : \mathcal{X} \times \mathcal{Y} \to \{0,1\}$

μ is a distribution over $\mathcal{X} \times \mathcal{Y}$.

$f_{\text{OR}}^k : \mathcal{X}^k \times \mathcal{Y} \to \{0,1\}$ is the problem of computing $f(x_1, y) \vee f(x_2, y) \vee \ldots \vee f(x_k, y)$ in the coordinator model, where P_i has input $x_i \in \mathcal{X}$ for each $i \in [k]$, and the coordinator has $y \in \mathcal{Y}$.

ν is a distribution on $\mathcal{X}^k \times \mathcal{Y}$: First pick $(X_1, Y) \sim \mu$, and then pick $X_2, \ldots, X_k$ from the conditional distribution $\mu \mid Y$.

Theorem (direct-sum). For any $f : \mathcal{X} \times \mathcal{Y} \to \{0,1\}$ and any distribution μ on $\mathcal{X} \times \mathcal{Y}$ for which $\mu(f^{-1}(1)) \leq 1/k^2$, we have $D_\nu^{1/k^3}(f_{\text{OR}}^k) = \Omega(k \cdot \text{ED}_\mu^{1/(100k^2)}(f))$. (Proof given later.)

A meta-problem: THRESH

In the THRESH_θ^n problem, site P_i ($i \in [k]$) holds an n-bit vector $x_i = \{x_{i,1}, \ldots, x_{i,n}\}$, and the k sites want to compute

$$\text{THRESH}_\theta^n(x_1, \ldots, x_k) = \begin{cases} 0, & \text{if } \sum_{j\in[n]}(\vee_{i\in[k]} x_{i,j}) \leq \theta, \\ 1, & \text{if } \sum_{j\in[n]}(\vee_{i\in[k]} x_{i,j}) \geq \theta + 1. \end{cases}$$

A meta-problem: THRESH

In the THRESH_θ^n problem, site P_i ($i \in [k]$) holds an n-bit vector $x_i = \{x_{i,1}, \ldots, x_{i,n}\}$, and the k sites want to compute

$$\text{THRESH}_\theta^n(x_1, \ldots, x_k) = \begin{cases} 0, & \text{if } \sum_{j\in[n]}(\vee_{i\in[k]} x_{i,j}) \le \theta, \\ 1, & \text{if } \sum_{j\in[n]}(\vee_{i\in[k]} x_{i,j}) \ge \theta + 1. \end{cases}$$

Theorem $\exists$ a distribution ζ, $D_\zeta^{1/k^4}(\text{THRESH}_\theta^n) = \Omega(kn)$.

Corollary $R^{1/3}(\text{THRESH}_\theta^n) = \Omega(kn/\log k)$

A meta-problem: THRESH

In the THRESH_θ^n problem, site P_i ($i \in [k]$) holds an n-bit vector $x_i = \{x_{i,1}, \ldots, x_{i,n}\}$, and the k sites want to compute

$$\text{THRESH}_\theta^n(x_1, \ldots, x_k) = \begin{cases} 0, & \text{if } \sum_{j \in [n]} (\vee_{i \in [k]} x_{i,j}) \leq \theta, \\ 1, & \text{if } \sum_{j \in [n]} (\vee_{i \in [k]} x_{i,j}) \geq \theta + 1. \end{cases}$$

Theorem $\exists$ a distribution ζ, $D_\zeta^{1/k^4}(\text{THRESH}_\theta^n) = \Omega(kn)$.
Corollary $R^{1/3}(\text{THRESH}_\theta^n) = \Omega(kn/\log k)$

The proof framework:

1. Choose f to be *2-DISJ* with an input distribution μ, and let *OR-DISJ*, ν be the corresponding f_{OR}^k and its input distribution. Apply direct-sum: $D_\nu^{1/k^3}(\text{OR-DISJ}) = \Omega(k \cdot ED_\mu^{1/(100k^2)}(\text{2-DISJ}))$.
2. Show that with very high probability, $\text{OR-DISJ}(X_1, \ldots, X_k, Y) = \text{THRESH}_\theta^n(X_1, \ldots, X_k)$ for some θ.

2-DISJ and OR-DISJ

$X \cap Y = \emptyset?$

$X \subseteq \{1, \ldots, n\}$ $Y \subseteq \{1, \ldots, n\}$

Exists a hard distribution τ_β, under which
$|X \cap Y| = 1$ (YES instance) w.p. β and
$|X \cap Y| = 0$ (NO instance) w.p. $1 - \beta$.

Theorem. (Generalization of [Razborov '90, BJKS '04])
$ED_{\tau_\beta}^{\beta/100}(\text{2-DISJ}) = \Omega(n)$

2-DISJ and OR-DISJ

$X \cap Y = \emptyset$?

$X \subseteq \{1, \ldots, n\}$ $\qquad$ $Y \subseteq \{1, \ldots, n\}$

Exists a hard distribution τ_β, under which
$|X \cap Y| = 1$ (YES instance) w.p. β and
$|X \cap Y| = 0$ (NO instance) w.p. $1 - \beta$.

Theorem. (Generalization of [Razborov '90, BJKS '04])
$ED_{\tau_\beta}^{\beta/100}(\text{2-DISJ}) = \Omega(n)$

By the direct-sum theorem, setting $\mu = \tau_\beta$ with $\beta = 1/k^2$, and ν constructed from μ, we have
$D_\nu^{1/k^3}(\text{OR-DISJ}) = \Omega(k \cdot ED_{\tau_\beta}^{\beta/100}(\text{2-DISJ})) = \Omega(kn)$.

Example: A reduction from THRESH to Connectivity

Reduction: An input $(X_1, \ldots, X_k)$ for THRESH $\Rightarrow$ A graph.

Each P_i creates an edge (u_i, v_j) for each $X_{i,j} = 1$. In addition, the coordinator reconstructs Y, and then creates a path containing $\{v_j \mid j \in Y\}$ and a path containing $\{v_j \mid j \in [r] \backslash Y\}$.

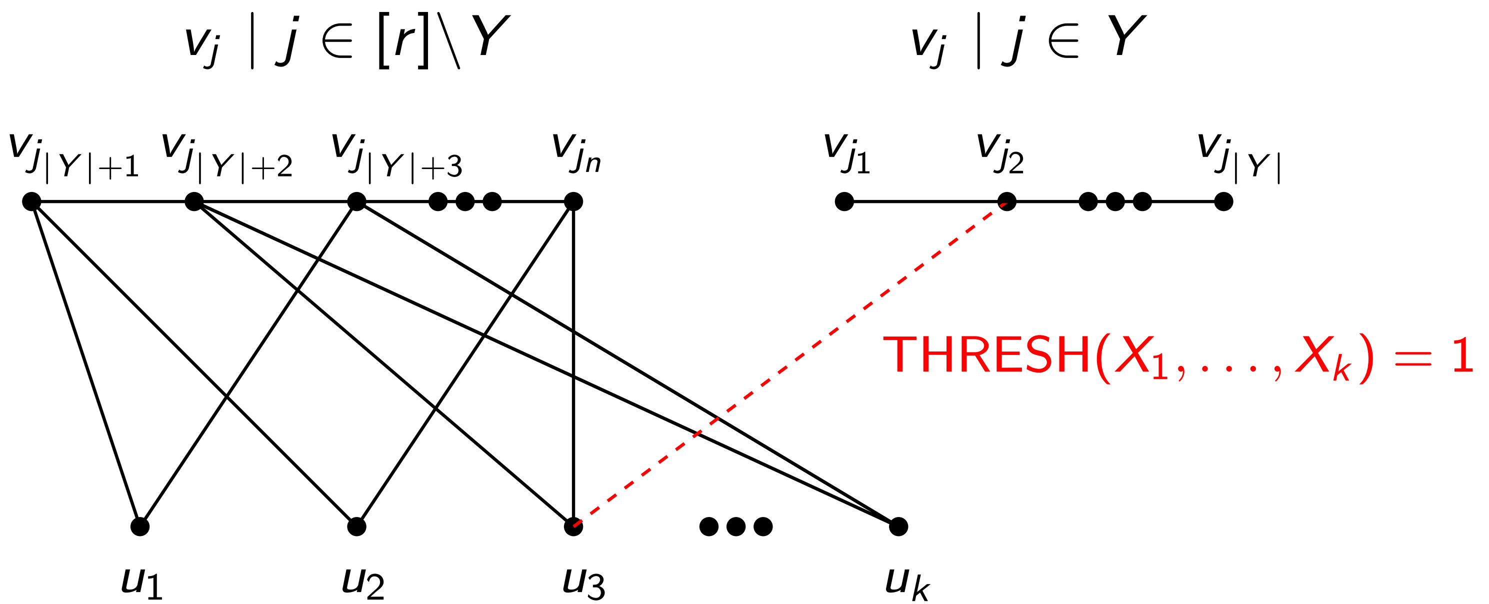

(u_i, v_j) exists (the graph is connected) if and only if $X_{i,j} = 1$

Proof of the direct-sum theorem

Proof of the direct-sum theorem

Theorem (direct-sum). For any $f : \mathcal{X} \times \mathcal{Y} \to \{0,1\}$ and any distribution μ on $\mathcal{X} \times \mathcal{Y}$ for which $\mu(f^{-1}(1)) \leq 1/k^2$, we have $D_{\nu}^{1/k^3}(f_{\mathrm{OR}}^k) = \Omega(k \cdot \mathrm{ED}_{\mu}^{1/(100k^2)}(f))$.

Proof of the direct-sum theorem

Theorem (direct-sum). For any $f : \mathcal{X} \times \mathcal{Y} \to \{0,1\}$ and any distribution μ on $\mathcal{X} \times \mathcal{Y}$ for which $\mu(f^{-1}(1)) \le 1/k^2$, we have $D_\nu^{1/k^3}(f_{\mathrm{OR}}^k) = \Omega(k \cdot \mathrm{ED}_\mu^{1/(100k^2)}(f))$.

The **proof** is by a reduction from a 2-player problem to a k-site problem.

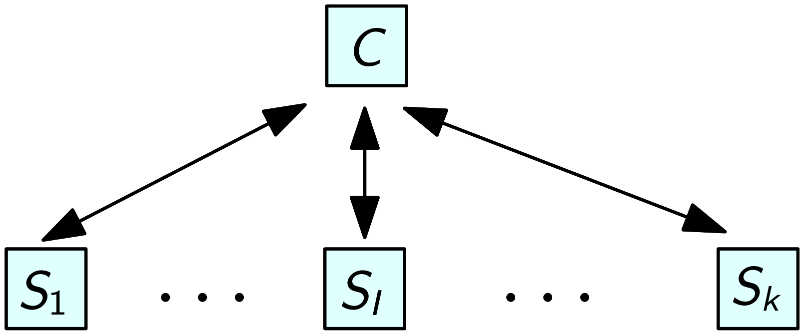

Proof of the direct-sum theorem

Theorem (direct-sum). For any $f : \mathcal{X} \times \mathcal{Y} \to \{0, 1\}$ and any distribution μ on $\mathcal{X} \times \mathcal{Y}$ for which $\mu(f^{-1}(1)) \leq 1/k^2$, we have $D_\nu^{1/k^3}(f_{\mathrm{OR}}^k) = \Omega(k \cdot \mathrm{ED}_\mu^{1/(100k^2)}(f))$.

The **proof** is by a reduction from a 2-player problem to a k-site problem.

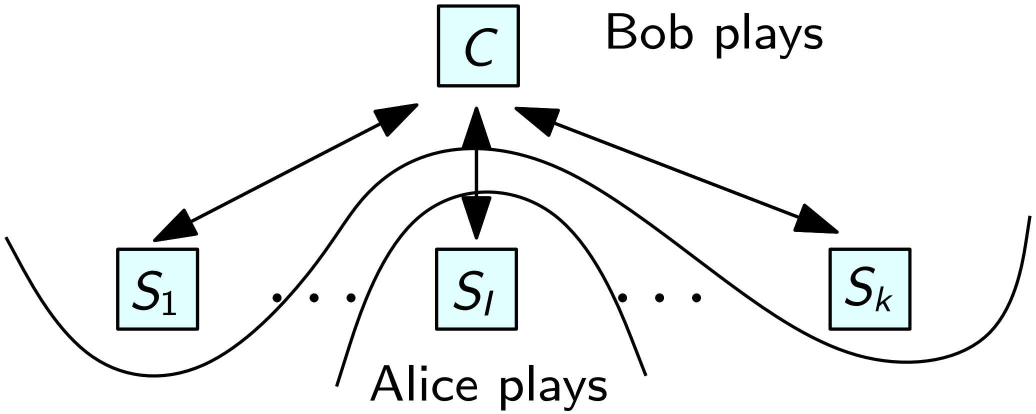

1. Alice and Bob have input $(X, Y) \sim \mu$ $(\mu(f^{-1}(1)) \leq 1/k^2)$

2. **Input reduction**: Alice picks a random site S_I and assigns it with input $X_I = X$. Bob plays the coordinator C and the rest $k-1$ sites. He assigns C with input Y, and S_i $(i \neq I)$ with input $X_i \sim \mu|Y$.

Proof of the direct-sum theorem

Theorem (direct-sum). For any $f : \mathcal{X} \times \mathcal{Y} \to \{0,1\}$ and any distribution μ on $\mathcal{X} \times \mathcal{Y}$ for which $\mu(f^{-1}(1)) \le 1/k^2$, we have $D_\nu^{1/k^3}(f_{\mathrm{OR}}^k) = \Omega(k \cdot \mathrm{ED}_\mu^{1/(100k^2)}(f))$.

The **proof** is by a reduction from a 2-player problem to a k-site problem.

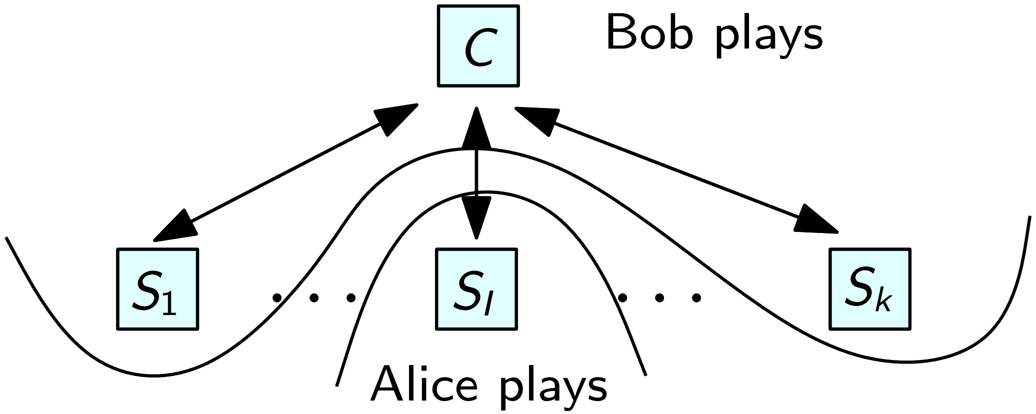

1. Alice and Bob have input $(X, Y) \sim \mu$ $(\mu(f^{-1}(1)) \le 1/k^2)$

2. **Input reduction**: Alice picks a random site S_I and assigns it with input $X_I = X$. Bob plays the coordinator C and the rest $k-1$ sites. He assigns C with input Y, and S_i $(i \neq I)$ with input $X_i \sim \mu | Y$.

3. They run a protocol for f_{OR}^k, w.pr. $1 - \frac{1}{k}$, $f(X_i, Y) = 0$ for all $i \neq I$, thus $f_{\mathrm{OR}}^k(X_1, \ldots, X_k, Y) = f(X_1, Y) \vee \ldots \vee f(X_k, Y) = f(X_I, Y) = f(X, Y)$.

Proof of the direct-sum theorem (cont.)

1. Alice and Bob have input $(X, Y) \sim \mu$ ($\mu(f^{-1}(1)) \leq 1/k^2$)

2. **Input reduction**: Alice picks a random site S_I and assigns it with input $X_I = X$. Bob plays the coordinator C and the rest $k-1$ sites. He assigns C with input Y, and S_i ($i \neq I$) with input $X_i \sim \mu | Y$.

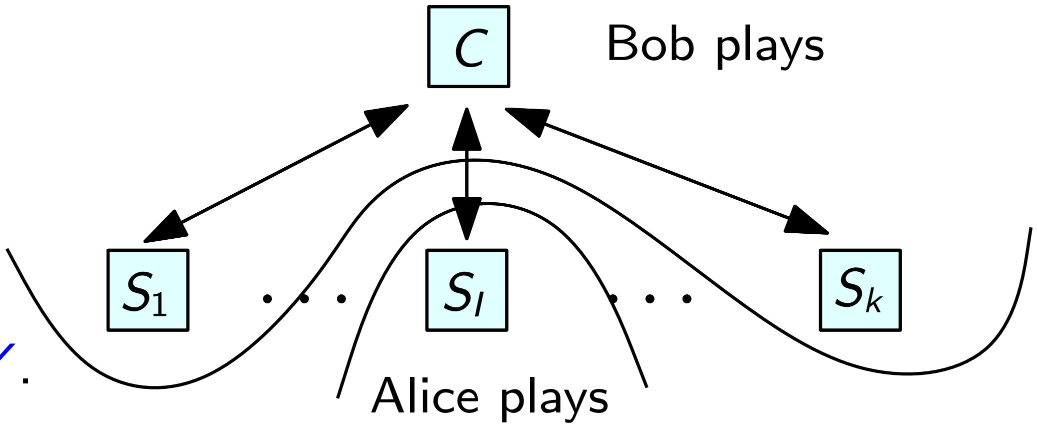

3. They run a protocol for f_{OR}^k, w.pr. $1 - \frac{1}{k}$, $f(X_i, Y) = 0$ for all $i \neq I$, thus $f_{\text{OR}}^k(X_1, \ldots, X_k, Y) = f(X_1, Y) \vee \ldots \vee f(X_k, Y) = f(X_I, Y) = f(X, Y)$.

4. They repeat the input reduction 3 times (using the same (X, Y)) and run the protocol for f_{OR}^k on each input, the probability that at least in one run (which Bob knows), $f(X_i, Y) = 0$ for all $i \neq I$, is $1 - 1/k^3$. Together with the fact that the error prob. of each run is at most $1/k^3$, the theorem follows from a union bound.

Proof of the direct-sum theorem (cont.)

1. Alice and Bob have input $(X, Y) \sim \mu$ ($\mu(f^{-1}(1)) \leq 1/k^2$)

$\mathrm{E}[\mathrm{CC}(\text{Alice, Bob})]$
$= 1/k \cdot \mathrm{CC}(k\text{-site problem})$

2. **Input reduction**: Alice picks a random site S_I and assigns it with input $X_I = X$. Bob plays the coordinator C and the rest $k-1$ sites. He assigns C with input Y, and S_i ($i \neq I$) with input $X_i \sim \mu | Y$.

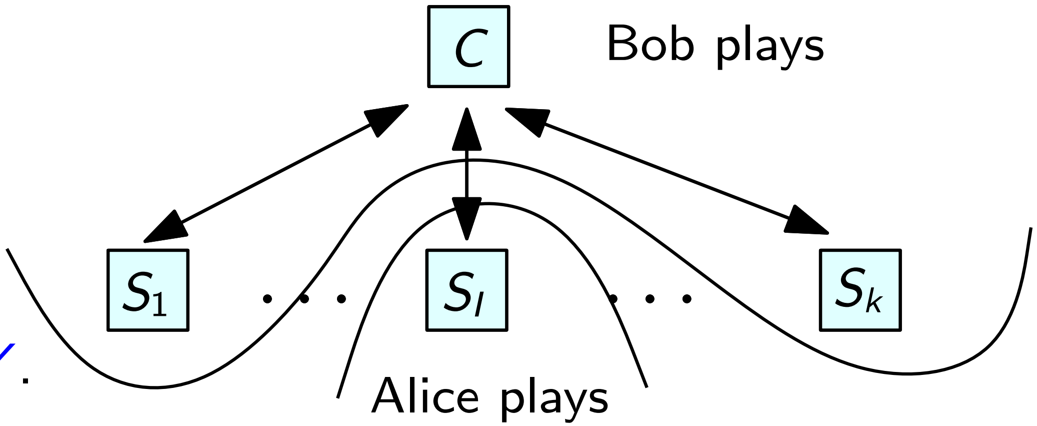

3. They run a protocol for f_{OR}^k, w.pr. $1 - \frac{1}{k}$, $f(X_i, Y) = 0$ for all $i \neq I$, thus $f_{\mathrm{OR}}^k(X_1, \ldots, X_k, Y) = f(X_1, Y) \vee \ldots \vee f(X_k, Y) = f(X_I, Y) = f(X, Y)$.

4. They repeat the input reduction 3 times (using the same (X, Y)) and run the protocol for f_{OR}^k on each input, the probability that at least in one run (which Bob knows), $f(X_i, Y) = 0$ for all $i \neq I$, is $1 - 1/k^3$. Together with the fact that the error prob. of each run is at most $1/k^3$, the theorem follows from a union bound.

Conclusions and open problems

- We attempt to deliver the following messages for protocol design in the coordinator model:
  - Allow an approximate solution.
  - Use well-designed input layouts.
  - Explore prior distributional properties of the input dataset.

Conclusions and open problems

- We attempt to deliver the following messages for protocol design in the coordinator model:
  - Allow an approximate solution.
  - Use well-designed input layouts.
  - Explore prior distributional properties of the input dataset.

- Can we obtain round-efficient protocols that (almost) match the lower bounds which hold even for round-inefficient protocols?

  Our protocols for bipartiteness and (approximate) diameter require $\Omega(\Delta)$ rounds where Δ is the diameter of the graph.

Thank you!
Questions?