On the Hardness of Information-Theoretic MPC

Yuval Ishai
Eyal Kushilevitz
Technion

Open Problems Museum

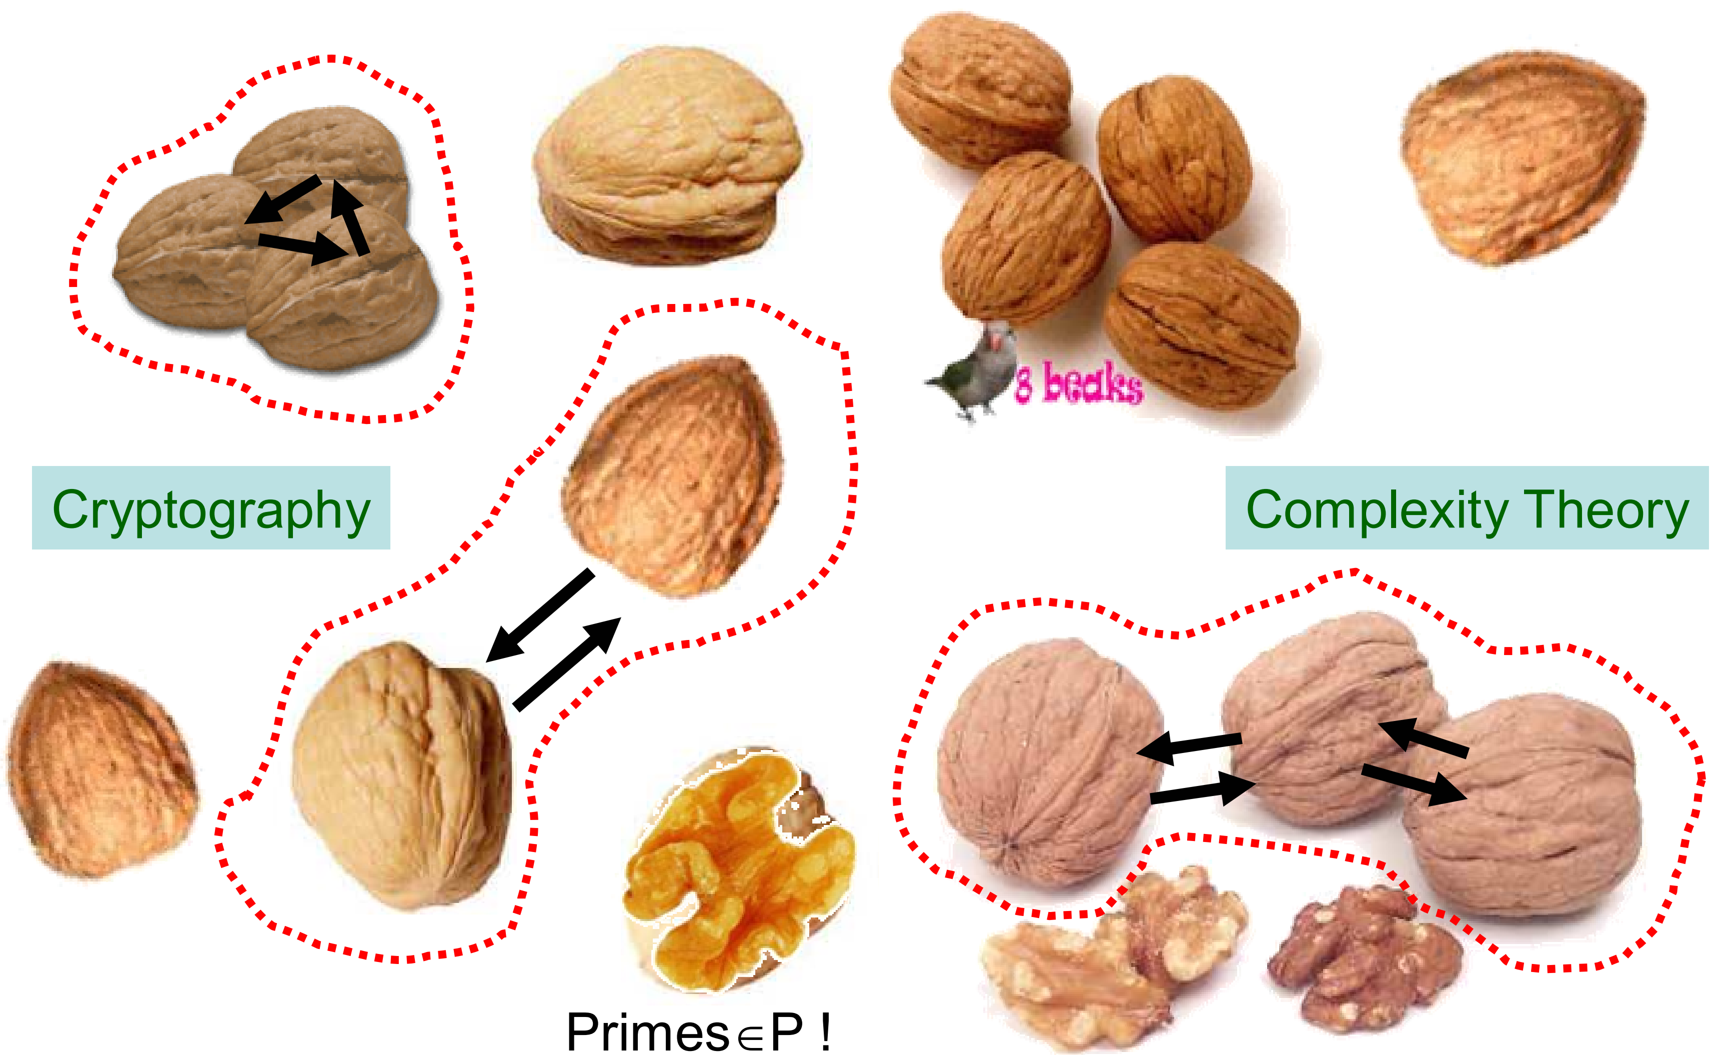

Motivating Question

Ben-Or, Goldwasser, Wigderson, 1988
Chaum, Crépeau, Damgård, 1988

Information-theoretic MPC is feasible!

$k \geq 3$ players can compute any function f of their inputs with total work = poly(circuit-size)

… or with work = poly(formula-size) and constant rounds [BB89,…]

Beaver, Micali, Rogaway, 1990
B., Feigenbaum, Kilian, R., 1990

Open question:

Can *k* computationally unbounded players compute an arbitrary f with communication = poly(input-length)?

Can this be done using a constant number of rounds?

Question Reformulated

Is the communication complexity of MPC strongly correlated with the computational complexity of the function being computed?

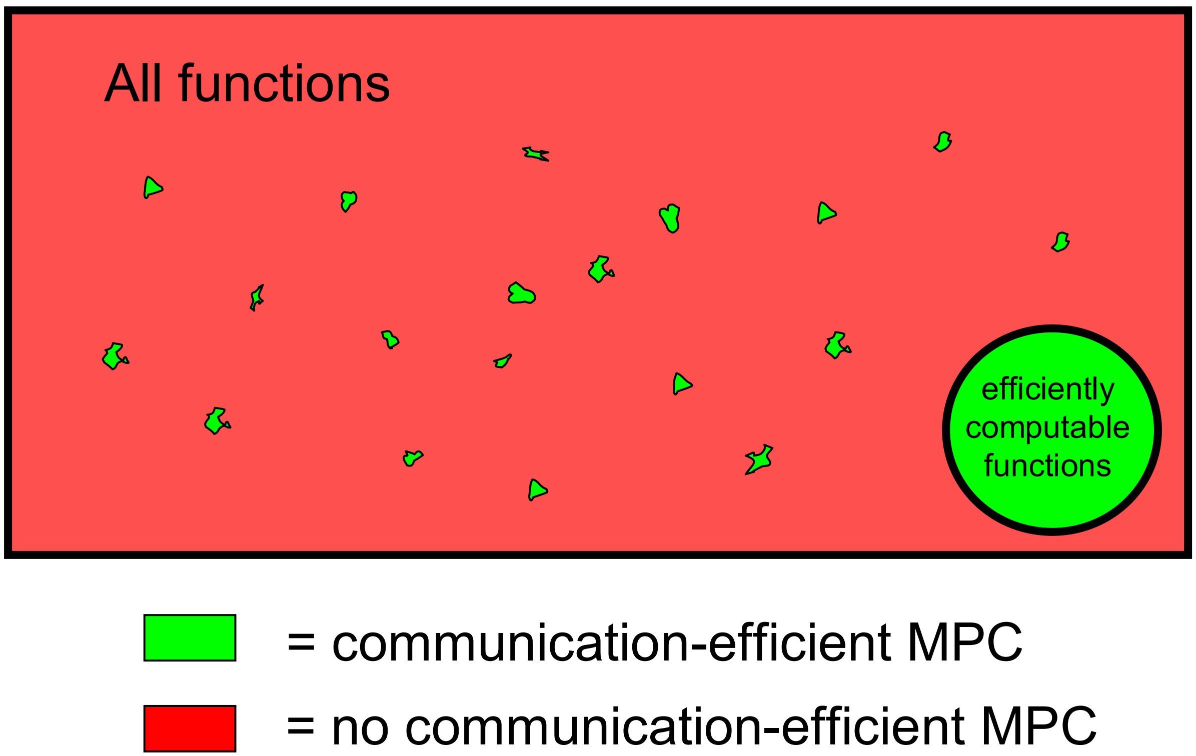

Our Results

- Connect latter MPC question to other, notoriously hard, problems.

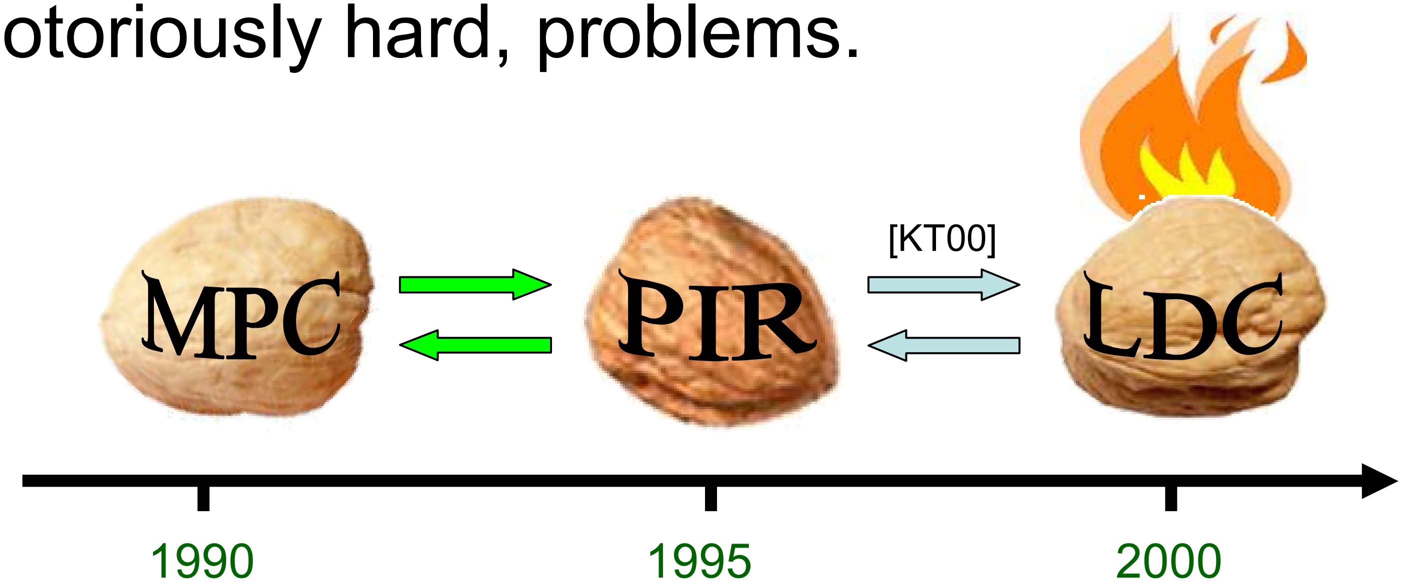

- The three problems are “essentially equivalent”
  – up to considerable deterioration of parameters

Significance

- Breakthrough on LDC question will imply breakthrough on MPC question and vice versa.
- Resolving MPC question is likely to be hard
  - Even when restricted to constant rounds

some f_0 cannot be computed by 18 unbounded players using polynomial communication and constant rounds

$\rightarrow$

breakthrough lower bound for LDC

Related Work

- Communication-efficient MPC with short inputs
  - communication=poly(input-length) possible when #players ≈ total input-length [BFKR90]
  - vacuous when #players << input-length.

- Communication-preserving secure computation [NN01]
  - Different goals: polynomial vs. sublinear communication
  - Different model: information-theoretic multi-party vs. computational two-party
  - Different techniques: our question is trivialized in the [NN01] model

Rest of Talk

- Describe primitives and questions:
  - MPC
  - PIR, SPIR
  - LDC
- Outline connections:

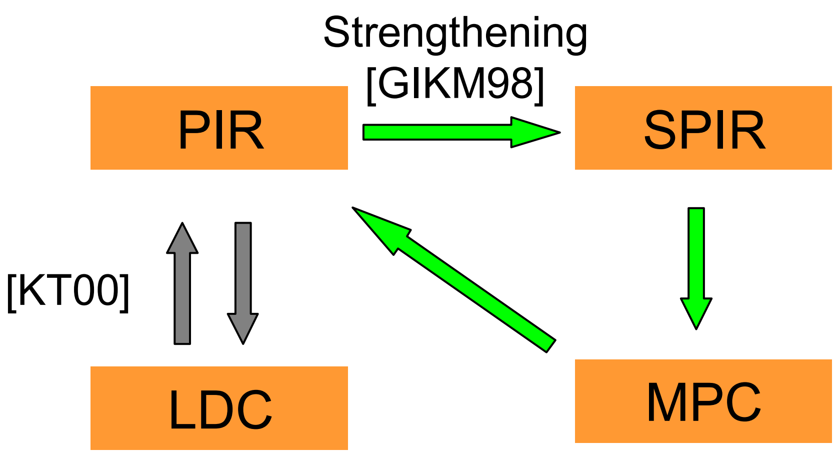

Secure Multiparty Computation (MPC)

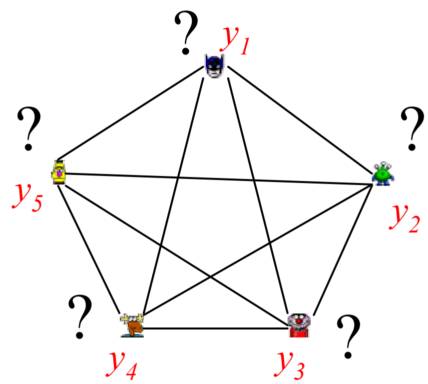

- f:$\{0,1\}^n \rightarrow \{0,1\}$
- k players
- Passive adversary
- Goal
  - Compute f
  - Maintain privacy against ≤ t players
- Parameters
  - k is fixed, input length n varies
  - This talk: t=1 (general t in addressed in paper)

Pure Information-Theoretic MPC

- Model
  - Secure channels
  - Computationally unbounded players
  - Security defined purely in terms of *information*
    - Compare to Shannon's notion of encryption
- Why is this model interesting?
  - Among "cleanest" nontrivial crypto problems!
  - Computation becomes feasible for small *n*
  - Useful for understanding the standard i.t. model
- Question: is there *k* for which every *f* can be computed using only poly(*n*) communication?

Information-Theoretic PIR [CGKS95]

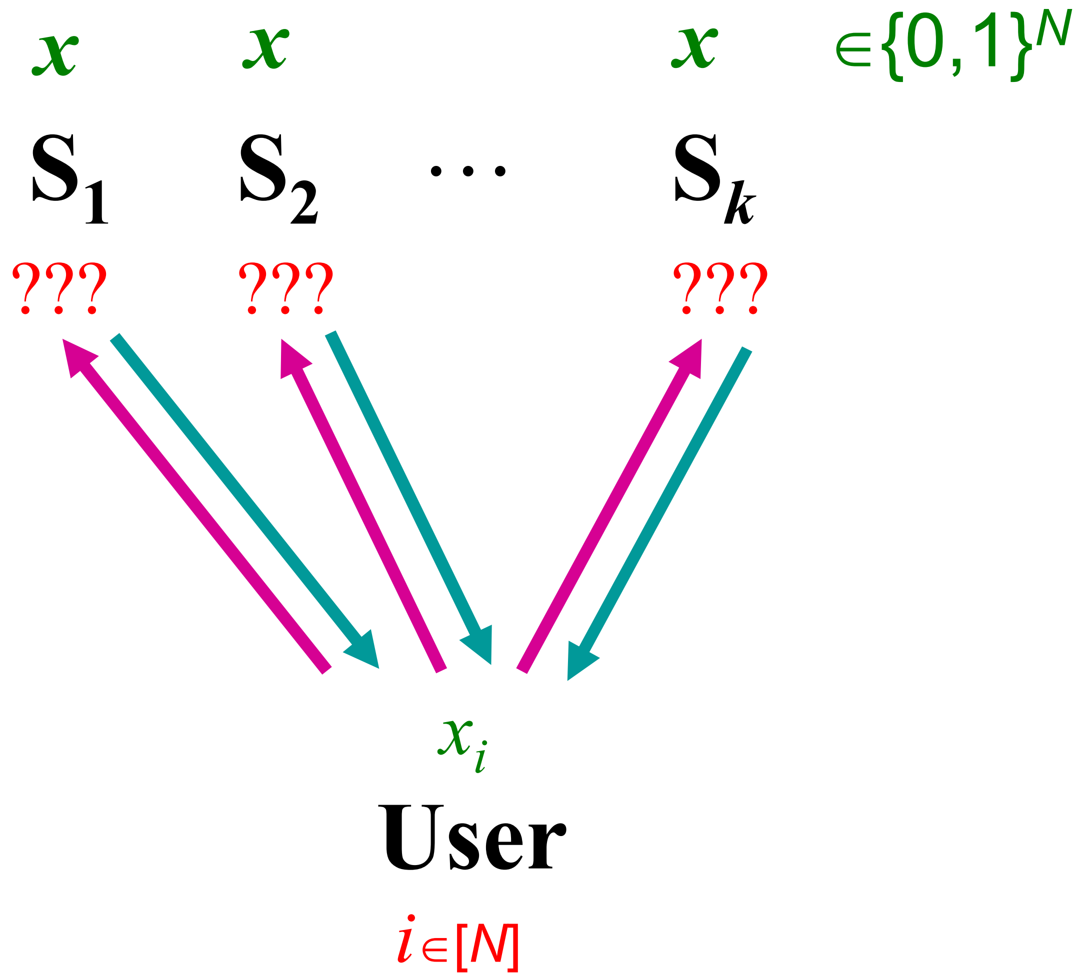

PIR (contd.)

- Goal: minimize communication complexity.
- Best upper bound: $N^{1/\tilde{\Omega}(k)}$ [CGKS95,…,BIKR02]
  - k=2: $O(N^{1/3})$
  - k=3: $O(N^{1/5.25})$
- Best lower bound: $\Omega(\log N)$ [Mann98]

- Question: polylog(N) communication?

Locally Decodable Codes (LDC) [KT00]

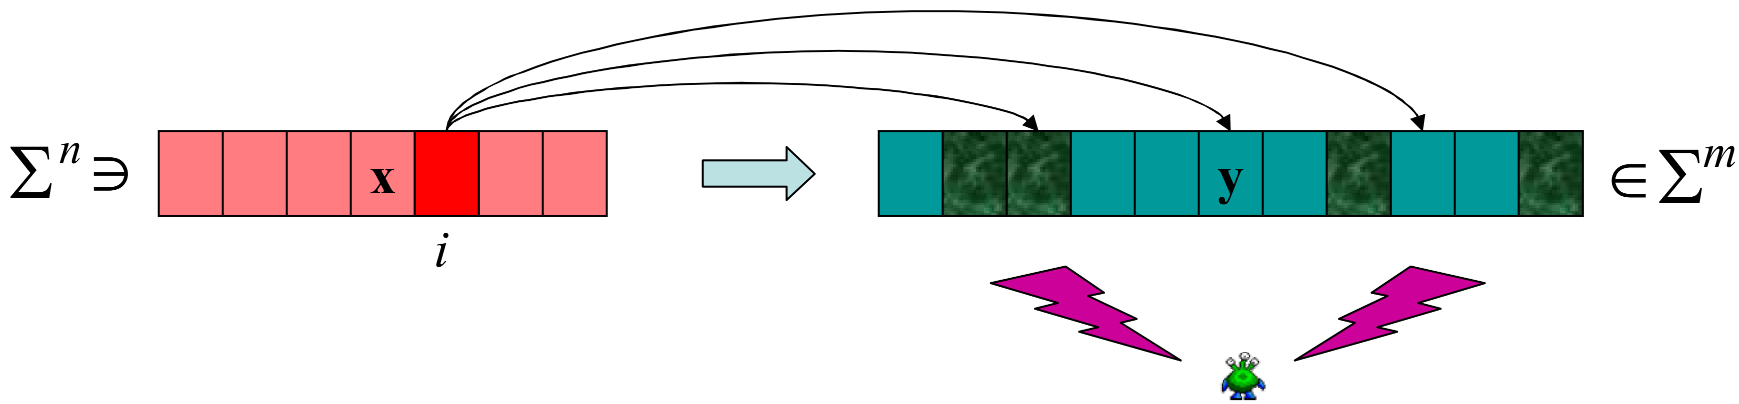

Requirements:

- High fault-tolerance
- Local decoding

	(q,δ,ε)-LDC	q-query LDC
High fault-tolerance	tolerate δm errors	const. $\delta,\varepsilon > 0$ (independent of n)
Local decoding	use q queries, succeed w/prob≥½+ε	

Question: Given q, how large should $m(n)$ be in a q-query LDC?

q=2: $2^{\Theta(n)}$ q=3: $2^{O(n^{0.5})}$ $\Omega(n^{1.5})$

PIR vs. LDC [KT00]

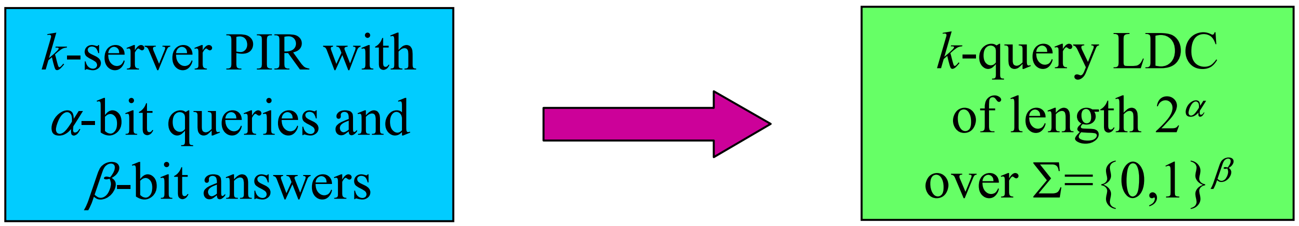

- Converse relation also holds.
- Best known LDC are obtained from PIR protocols.
  - const. q: $m=exp(n^{c\cdot \log\log q / q\log q})$
- k-server polylog PIR $\leftrightarrow$ k-query “quasi-poly” LDC

Symmetric PIR (SPIR) [GIKM98]

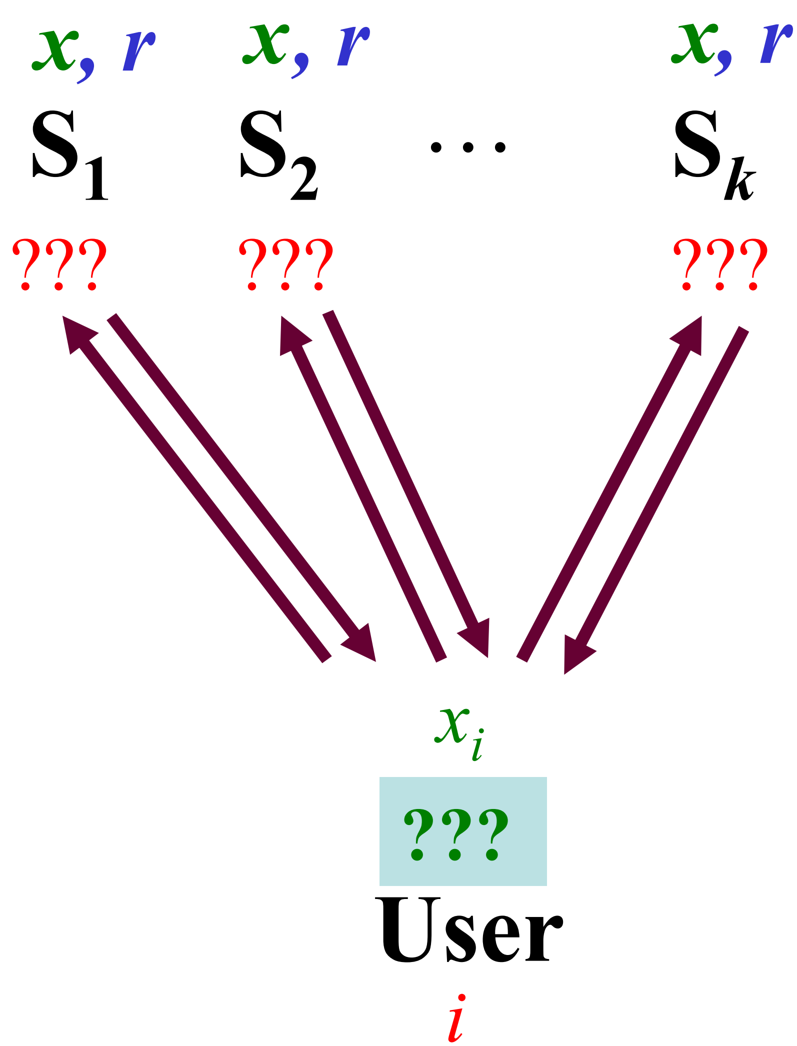

PIR→SPIR

- General PIR→SPIR transformation [GIKM98]
  - low communication overhead
  - one additional server
  - $|r|=N$: way too much for our purposes!
- Our approach: information-theoretic derandomization
  - Idea: if CC(SPIR)=c, then $\exists S \subseteq \{0,1\}^N$ of size $\approx 2^{c+\sigma}$ such that $r \in_R S$ is as good as $r \in_R \{0,1\}^N$, up to $2^{-\sigma}$ statistical distance.
  - ➢ SPIR protocols do not require much more randomness than communication.
  - Similar result can be shown for arbitrary i.t. protocols.

SPIR→MPC

k servers k^2+2 players

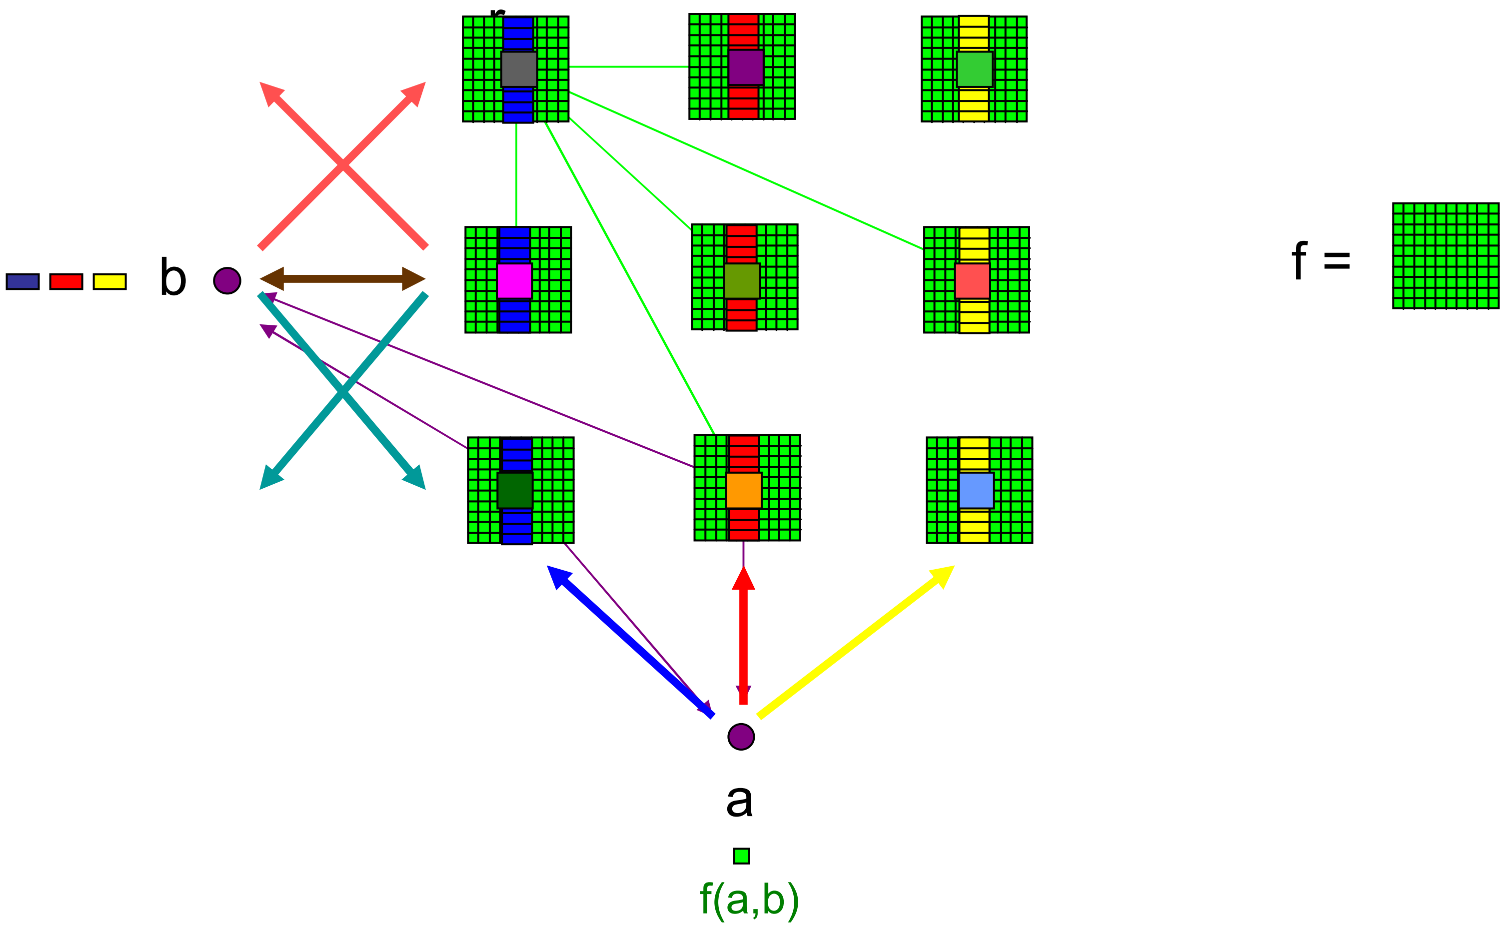

MPC→PIR

- Idea:
  - view database x as a truth-table of f_x
  - apply MPC among servers to let user privately learn $f_x(i)$
  - Some massaging required
- Produces multi-round PIR
- Still good enough to get nontrivial LDC

Further Research

- Find more connections
  - Generalized secret-sharing?
  - *Time* efficient constant-round protocol for any $f \in P$?
- Improve PIR→MPC connection
  - Multi-round PIR → MPC?
  - Eliminate growth of k?
- Computationally efficient derandomization
  - Easy given exponentially strong PRGs
  - Can one use standard PRGs?
  - Better yet, on worst-case hardness assumptions (a-la Nisan-Wigderson)?