Secure Multi-Party Differential Privacy

Peter Kairouz, Sewoong Oh, and Pramod Viswanath
E-mails: {kairouz2,swoh,pramodv}@illinois.edu
University of Illinois at Urbana Champaign, USA

Secure Multi-Party Differential Privacy

Private multi-party computation:

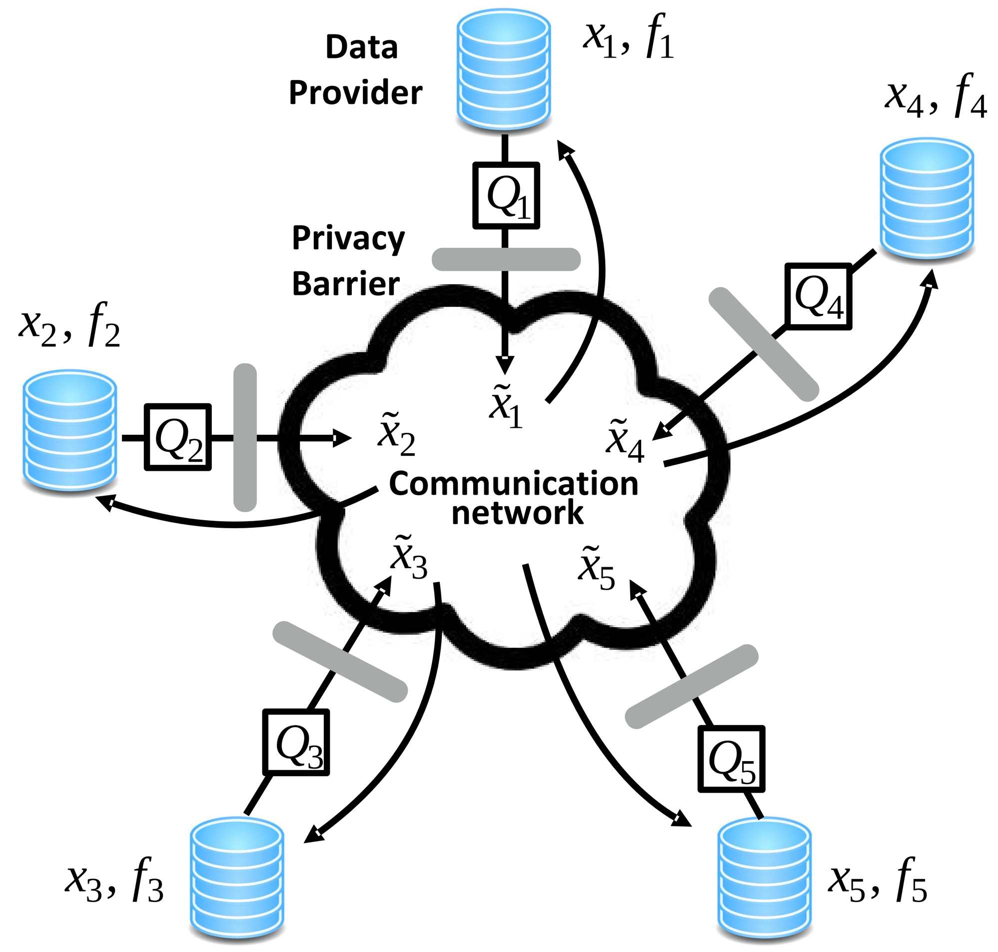

private multi-party computation

- parties exchange information to compute their functions
- central observer interested in computing a separate function f_0
- $x_1, x_2, \cdots, x_5$ are independent **binary** variables
- important setting in distributed in statistics and cloud computing

Interactive vs. non-interactive mechanisms:

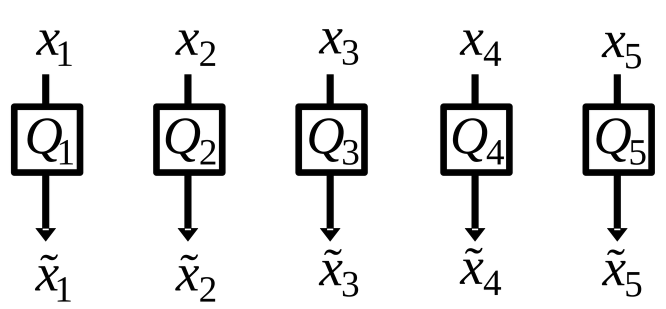

non-interactive mechanisms

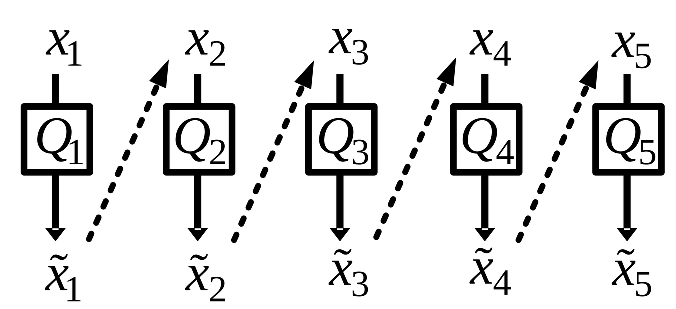

interactive mechanisms

- a more general representation:

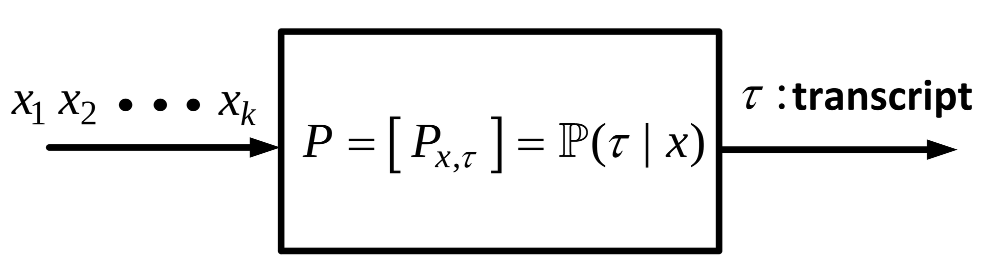

multi-party privatization mechanism $P_{x,\tilde{x}}$

- $P \in [0,1]^{2^k \times |\mathcal{T}|}$, where $\mathcal{T}$ is the space of all output transcripts
- $x = (x_1, x_2, \cdots, x_k)$

$$\tau\text{-th column of } P \text{ is a rank 1 tensor} \implies P(x|\tau) = \prod_i P(x_i|\tau)$$

Local differential privacy:

- A mechanism P is $\{e^{\varepsilon_i}, \delta_i\}$-differentially private if

$$\mathbb{P}(\tau|x_i, x_{-i}) \le e^{\varepsilon_i}\,\mathbb{P}(\tau|x_i', x_{-i}) + \delta_i \quad \forall i, x_i, x_i', x_{-i}, \tau$$

- $x_{-i} = (x_1, \cdots, x_{i-1}, x_{i+1}, \cdots, x_k)$

The Fundamental Privacy-Utility Tradeoff

Function estimation:

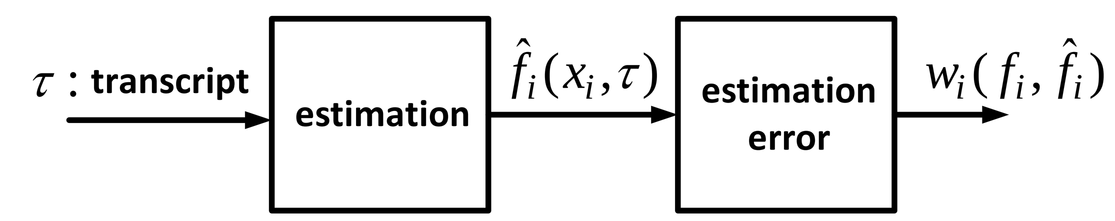

- user i estimates f_i using τ and x_i
- the central observer estimates f_0 using τ

Average accuracy case:

$$\mathrm{ACC}_{\mathrm{ave}}(P, w_i, f_i, \hat{f}_i) \equiv \frac{1}{2^k} \sum_{x \in \{0,1\}^k} \mathbb{E}_{\hat{f}_i, P_{x,\tau}}[w_i(f_i(x), \hat{f}_i(\tau, x_i))]$$

- if $w_i(y, y') = \mathbb{I}_{(y=y')}$ then $\mathrm{ACC}_{\mathrm{ave}}$ = probability of correct estimation
- for a fixed $P_{x,\tau}$, the optimal estimation rule is

$$\hat{f}_{i,\mathrm{opt}}(\tau, x_i) = \arg\max_y \sum_{x_{-i} \in \{0,1\}^{k-1}} P_{x,\tau}\, w_i(f_i(x), y)$$

Fundamental Privacy-Utility Tradeoff:

- maximize accuracy subject to privacy constraints

$$\begin{aligned} \underset{P, \hat{f}_i}{\text{maximize}} \quad & \mathrm{ACC}_{\mathrm{ave}}(P, w_i, f_i, \hat{f}_i), \\ \text{subject to} \quad & P \text{ and } \hat{f}_i \text{ are row-stochastic matrices},\ \mathrm{rank}(P^{(\tau)}) = 1 \quad \forall \tau \\ & P_{(x_i, x_{-i}),\tau} \le e^{\varepsilon_i} P_{(x_i', x_{-i}),\tau} + \delta_i \quad \forall i, x_i, x_i', x_{-i}, \tau \end{aligned}$$

- $P^{(\tau)}$ is the k-th order tensor of the τ-th column of P

The randomized response mechanism:

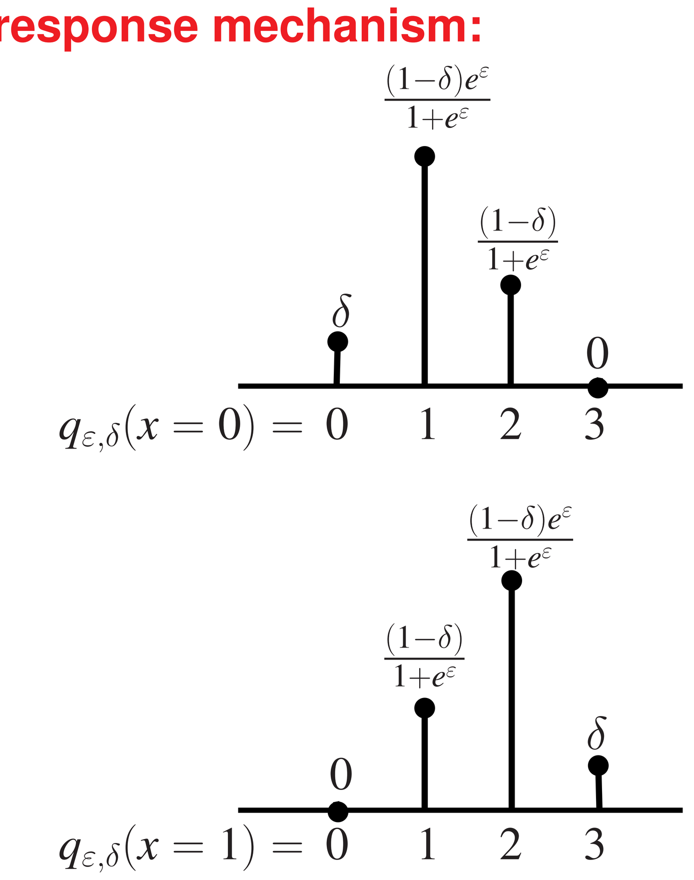

The Optimality of the Randomized Response Mechanism

For any pair $(e^{\varepsilon_i}, \delta_i)$, any function f_i, and any accuracy measure w_i, the randomized response, along with its corresponding optimal estimation rule, achieves the maximum accuracy for the i-th party, among all $\{e^{\varepsilon_i}, \delta_i\}$-differentially private interactive protocols and all estimation rules.

- **interaction** is not needed!
- Randomized response is also optimal for the **worst case accuracy**

Multi-Party XOR

Multi-Party XOR

Consider k-party computation for $f_0(x) = x_1 \oplus \cdots \oplus x_k$, and the estimation accuracy measure is one if correct and zero if not, i.e. $w_0(0,0) = w_0(1,1) = 1$ and $w_0(0,1) = w_0(1,0) = 0$. For any $\{e^{\varepsilon}, 0\}$-differentially private protocol P and any decision rule $\hat{f}$, the average case accuracy is bounded by

$$\mathrm{ACC}_{\mathrm{ave}}(P, w_0, f_0, \hat{f}_0) \le \frac{\sum_{i=0}^{\lfloor k/2 \rfloor} \binom{k}{2i} e^{\varepsilon(k-2i)}}{(1+e^{\varepsilon})^k},$$

where equality is achieved by the randomized response

- optimal estimation rule: XOR all the received privatized bits
- when $\varepsilon \simeq 0$, $\mathrm{ACC}_{\mathrm{ave}} = 0.5 + 2^{-(k+1)}\varepsilon^k + O(\varepsilon^{k+1})$

Generalization to Multiple Bits

- one party with one bit x and the second party has two bits y_1 and y_2

$$f(x, y_1, y_2) = \begin{cases} y_1 \oplus y_2 & \text{if } x = 0\,, \\ y_1 \wedge y_2 & \text{if } x = 1\,. \end{cases}$$

- randomized response: publish privatized versions of x, y_1, and y_2
- interactive mechanism: party 2 observes $\tilde{x}$ and privatizes
  - $y_1 \oplus y_2$ if $\tilde{x} = 0$
  - $y_1 \wedge y_2$ if $\tilde{x} = 1$
- estimation accuracy is measured by the Hamming distance

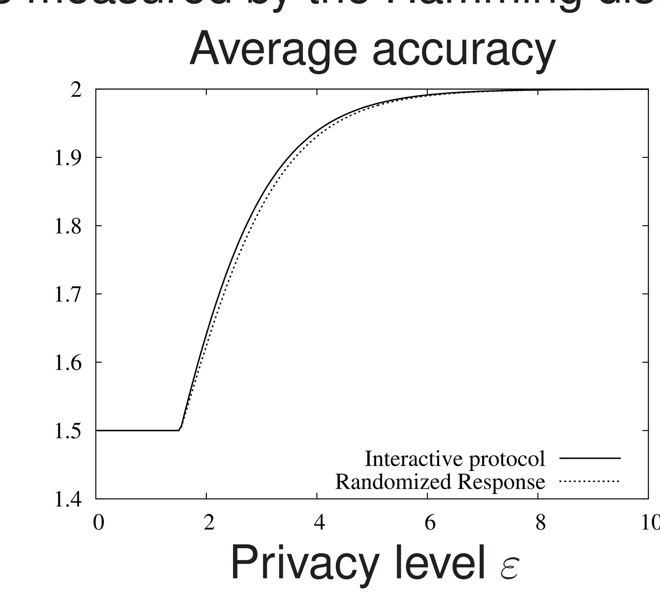

Going Forward

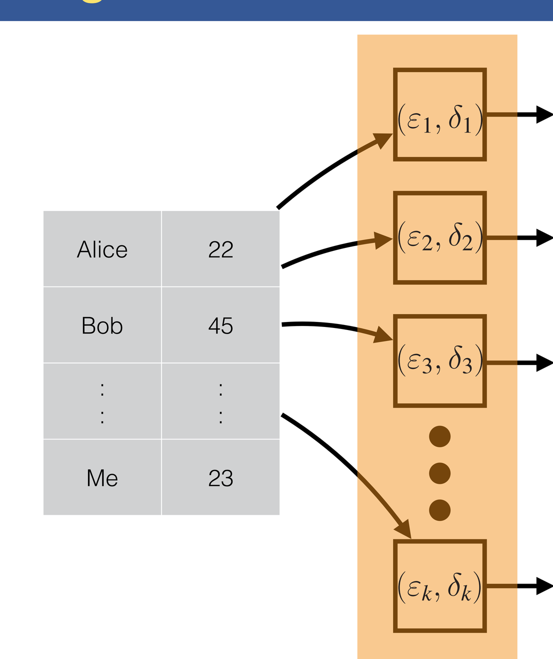

Composition Attacks

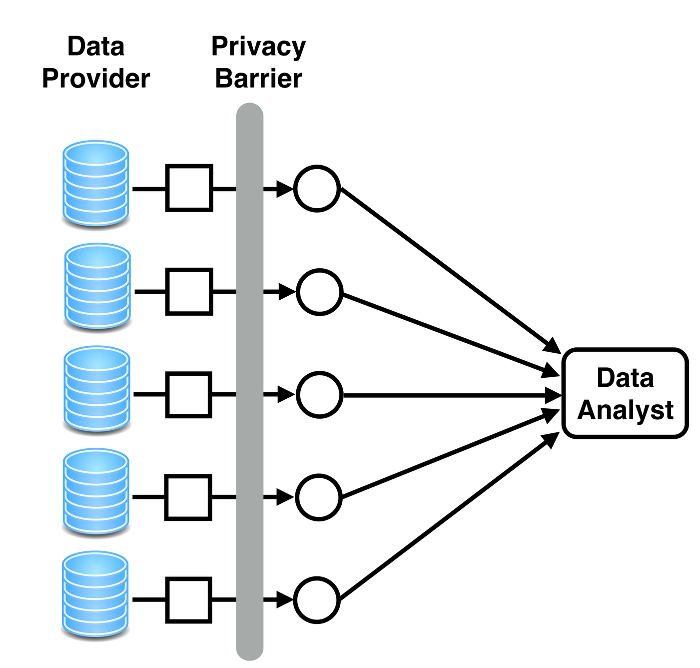

Local Privacy

- *"The Composition Theorem in Differential Privacy"*, ICML 2015
- *"Extremal Mechanisms for Local Differential Privacy"*, NIPS 2014