Discrete Distribution Estimation under Local Privacy

Peter Kairouz, Keith Bonawitz, and Daniel Ramage
e-mails: kairouz2@illinois.edu, bonawitz@google.com, and dramage@google.com

Distribution estimation

We need to understand **patterns across large groups** but **do not need to look at any individual**.

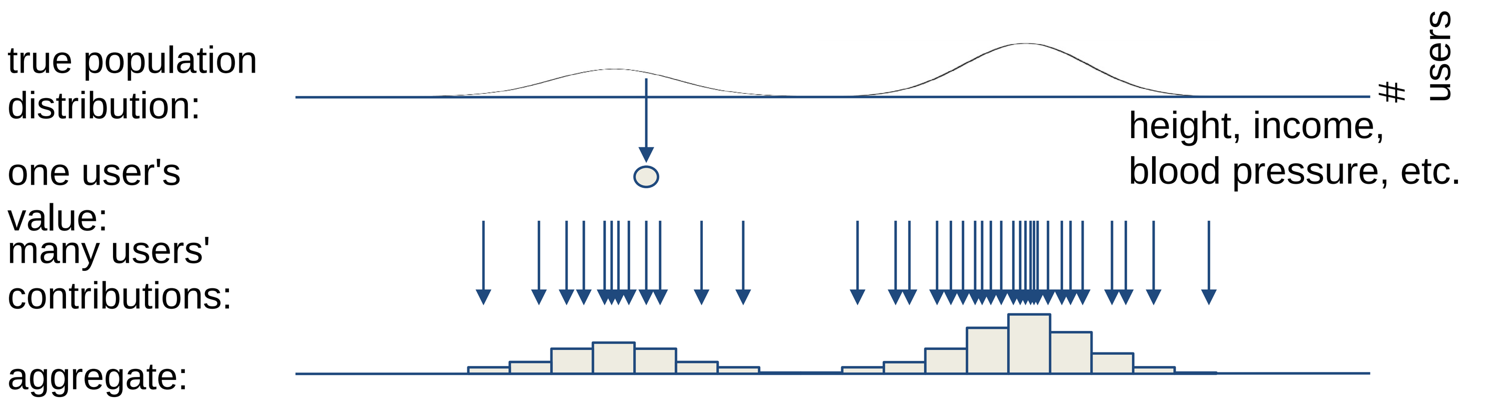

What's wrong with current approaches

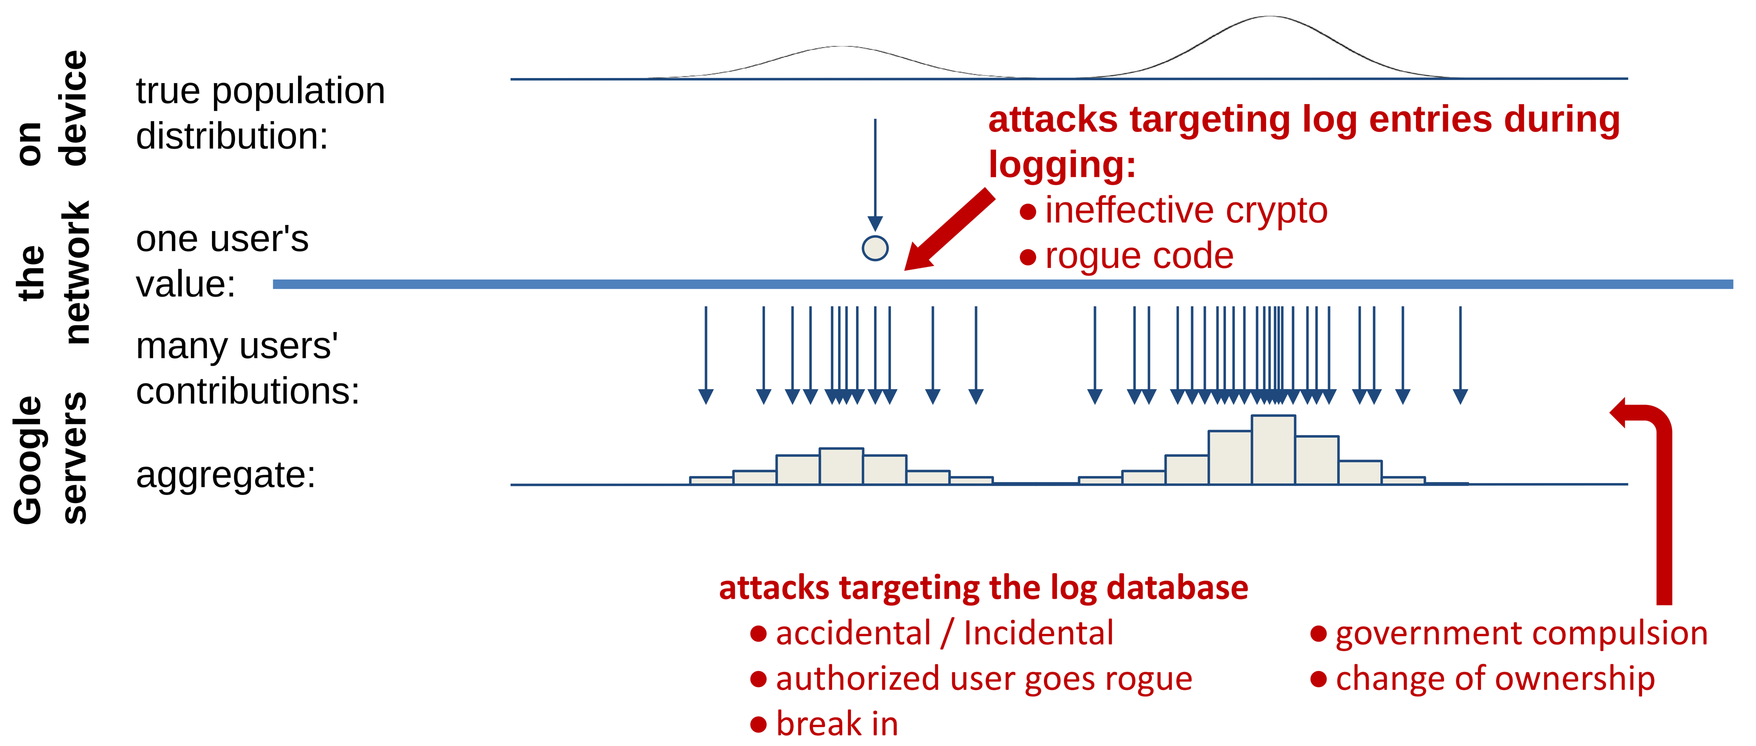

Private distribution estimation

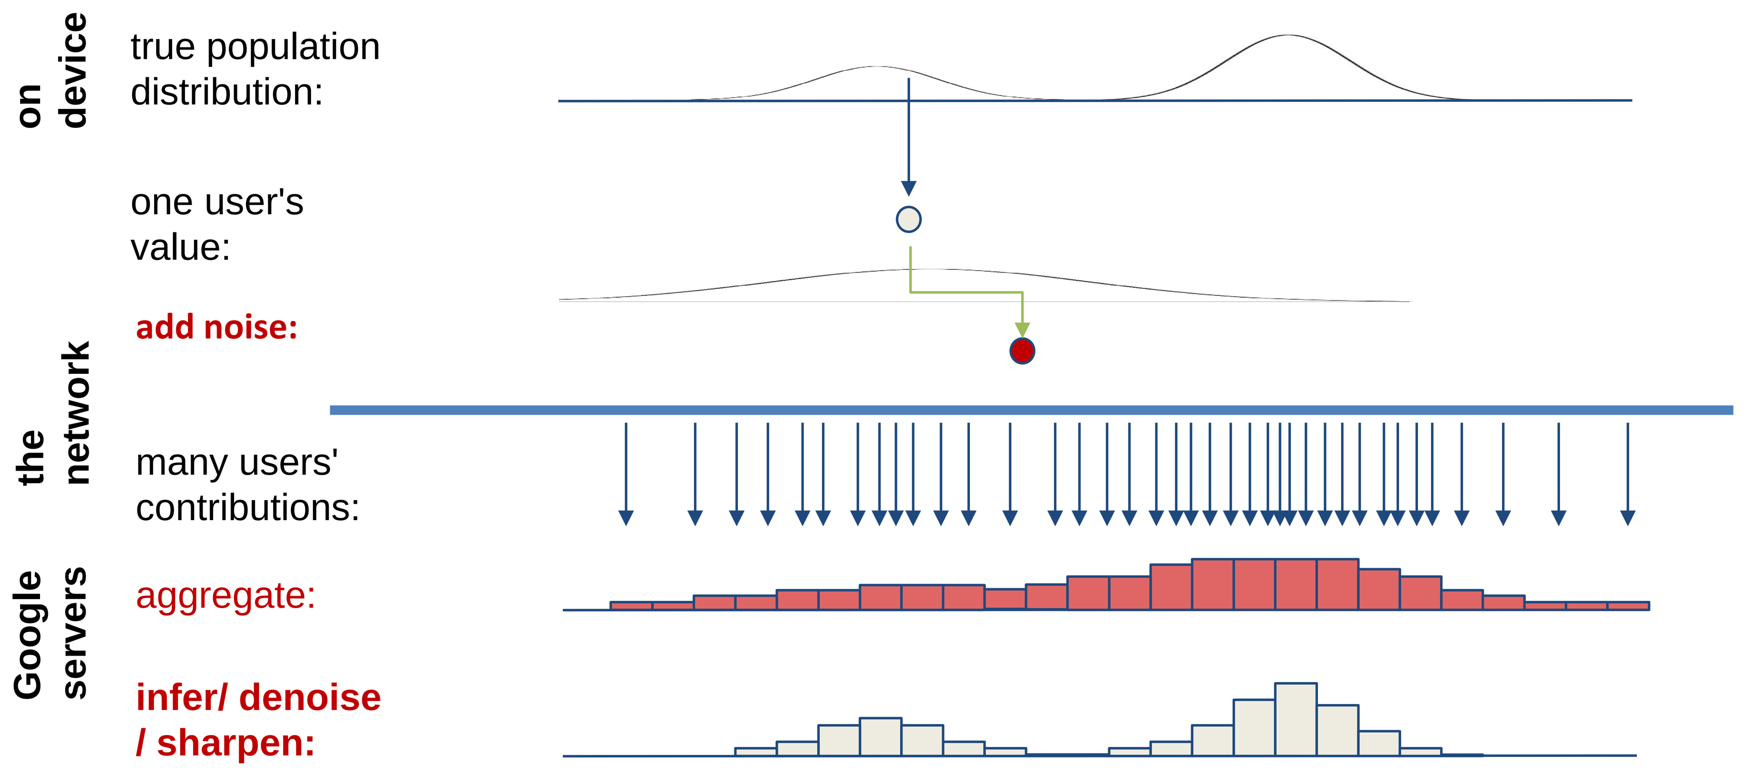

Less noise: less privacy but easier to denoise
More noise: more privacy but more data required to denoise

Local differential privacy

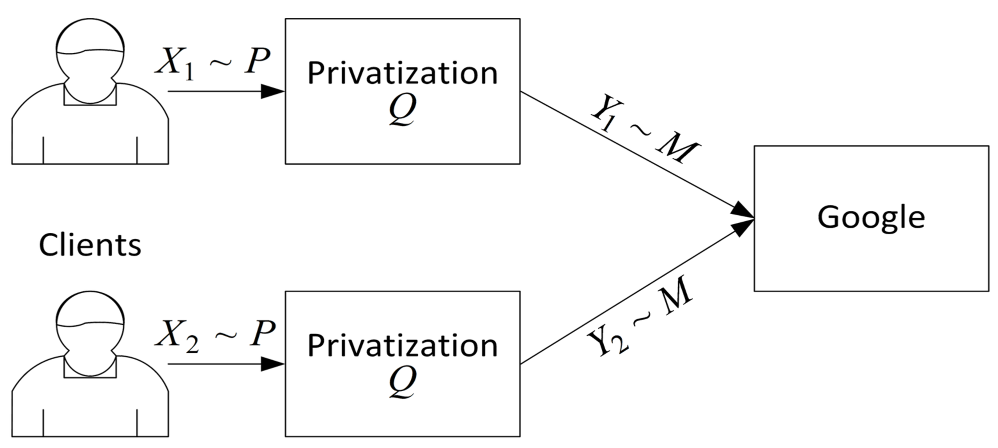

- If true answer is x, say y with probability: $Q(Y = y|X = x)$
- Q is locally differentially private if: $e^{-\varepsilon} \le \frac{Q(Y=y|X=x)}{Q(Y=y|X=x')} \le e^{\varepsilon} \quad \forall x, x', y$

Fundamental privacy-utility tradeoff

$$\inf_Q \inf_{\hat{P}} \sup_P \mathbb{E}\ \ell(P, \hat{P}(Q))$$

$$\text{subject to } Q \text{ locally differentially private}$$

- Why worst case over all dirbitions?
  - otherwise Q can be a trivial function of P
- Why is this problem a hard one?
  - Because minimax estimation without privacy is already hard

- What do we already know about this problem?

$$\inf_Q \inf_{\hat{P}} \sup_P \mathbb{E}\ \ell_2^2(P, \hat{P}(Q)) \approx \frac{k}{n\varepsilon^2}, \quad \text{for } \varepsilon \in [0,1]$$

$$\inf_Q \inf_{\hat{P}} \sup_P \mathbb{E}\ \ell_1(P, \hat{P}(Q)) \approx \frac{k}{\sqrt{n\varepsilon^2}}, \quad \text{for } \varepsilon \in [0,1]$$

What privacy mechanisms achieve the fundamental privacy-utility tradeoff for various privacy levels and alphabet sizes?

Binary alphabets: Warner's randomized response

Have you ever used illegal drugs?

answer **truthfully** w.p. $\frac{e^\varepsilon}{e^\varepsilon + 1}$ **lie** w.p. $\frac{1}{e^\varepsilon + 1}$

For all loss functions and all privacy levels, Warner's randomized response achieves the best privacy-utility tradeoff.

k-ary alphabets

- How do we generalize Warner's randomized response:
  1. modify the encoding: *k*-RR
  2. modify the mechanism: *k*-RAPPOR

k-ary Randomized Response (*k*-RR)

answer **truthfully** w.p. $\frac{e^\varepsilon}{e^\varepsilon + k - 1}$ **lie** w.p. $\frac{k-1}{e^\varepsilon + k - 1}$

k-ary Randomized Aggregatable Privacy Preserving Ordinal Response (*k*-RAPPOR)

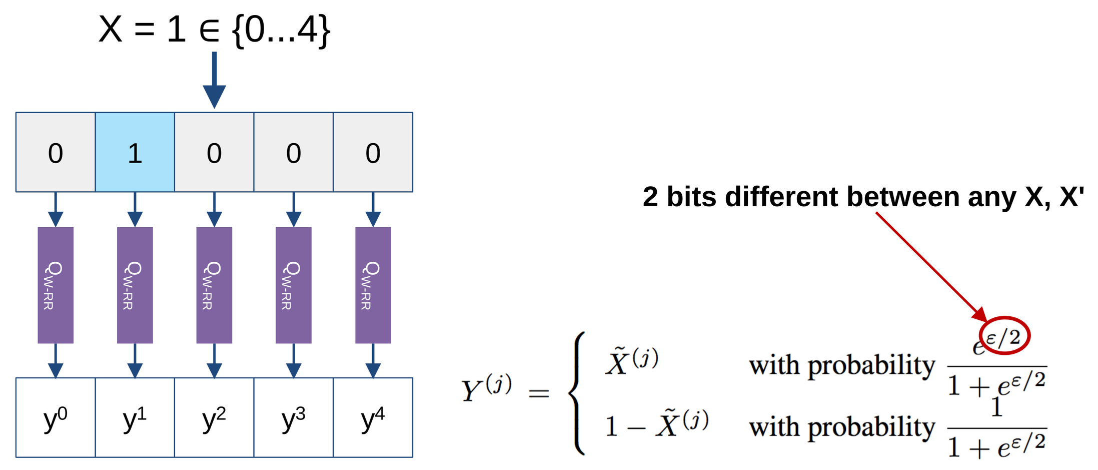

$$Y^{(j)} = \begin{cases} \tilde{X}^{(j)} & \text{with probability } \frac{e^{\varepsilon/2}}{1+e^{\varepsilon/2}} \\ 1 - \tilde{X}^{(j)} & \text{with probability } \frac{1}{1+e^{\varepsilon/2}} \end{cases}$$

- **For l1 and l2 loss functions, *k*-RR is order-optimal in the low privacy regime and strictly suboptimal in the high privacy regime**
- **For l1 and l2 loss functions, *k*-RAPPOR is order-optimal in the high privacy regime and strictly suboptimal in the low privacy regime**

- Sample complexity under both schemes

	ε ≈ ln(k) (Low Privacy)	**Small ε (High Privacy)**	**General**
No Privatization	n	n	n
***k*-RR**	$n/4$	$n\varepsilon^2/k^2$	$n\left(\frac{e^\varepsilon - 1}{e^\varepsilon + k - 1}\right)^2$
***k*-RAPPOR**	$n/\sqrt{k}$	$n\varepsilon^2/4k$	$n\left(\frac{(k-1)^2(e^{\varepsilon/2}-1)^2}{(k-1)(e^{\varepsilon/2}-1)^2 + k^2 e^\varepsilon}\right)$

Open alphabets

- What if we don't know the set of input symbols ahead of time?

Open Alphabets Randomized Response (O-RR)

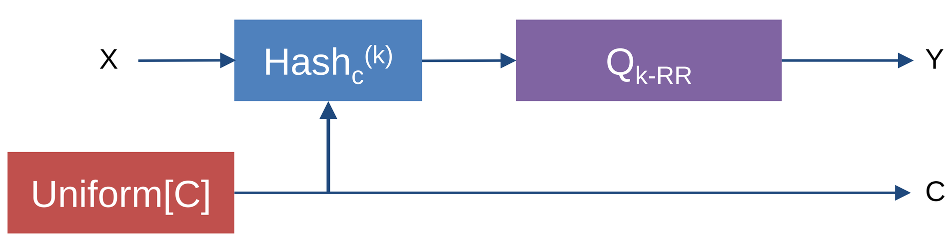

Open Alphabets RAPPOR (O-RAPPOR)

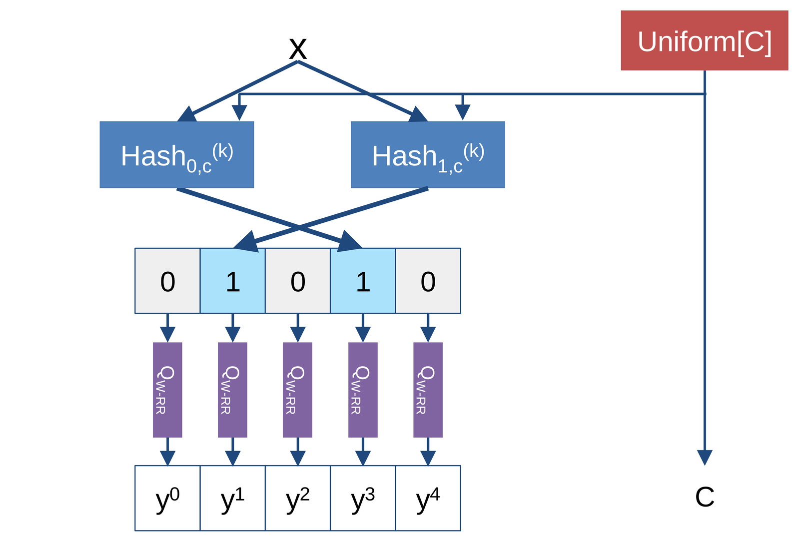

- Simulation results:

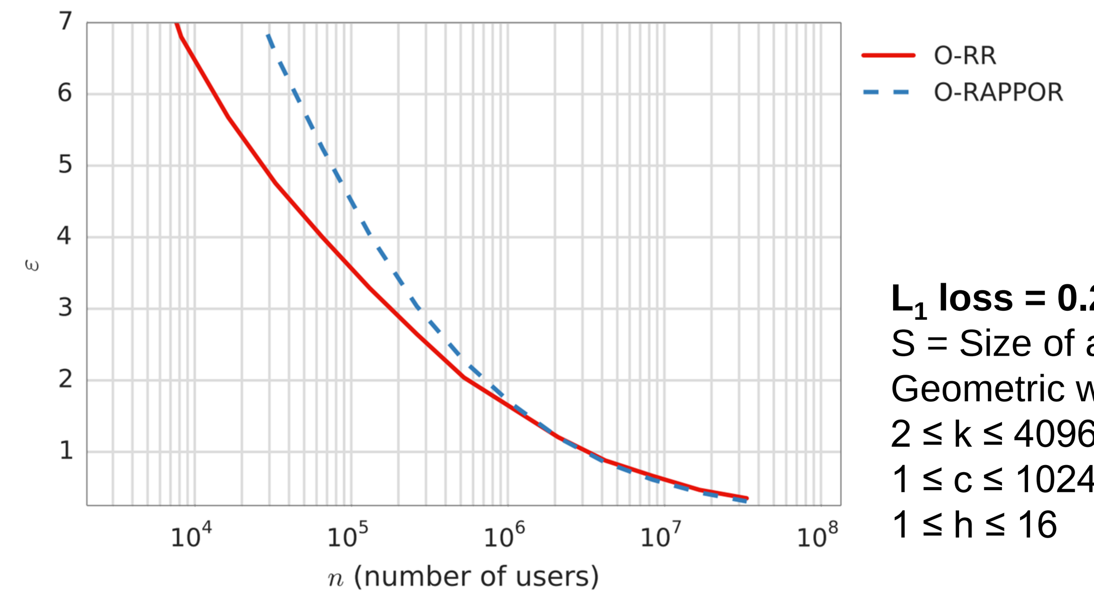

L_1 loss = 0.20;
S = Size of alphabet =256;
Geometric with mean= S/5
2 ≤ k ≤ 4096
1 ≤ c ≤ 1024
1 ≤ h ≤ 16

Closed alphabets: revisited

- A **Minimal Perfect Hash Function** maps m keys to m consecutive integers.
- **For Closed Sets:** Modify O-RR and O-RAPPOR to use Minimal Perfect Hash Functions.
- Note that with C=1 and h=1, we recover k-RR and k-RAPPOR

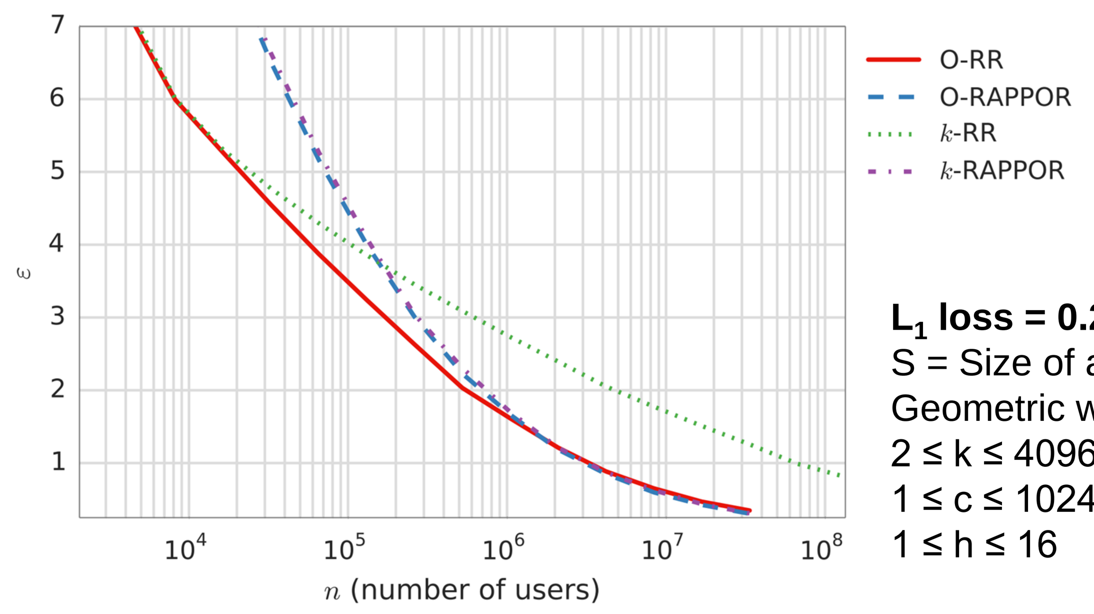

L_1 loss = 0.20;
S = Size of alphabet =256;
Geometric with mean=S/5
2 ≤ k ≤ 4096
1 ≤ c ≤ 1024
1 ≤ h ≤ 16

Acknowledgments

- Ulfar Erlingsson
- Ilya Mironov
- Andrey Zhmoginov