

# Practical Locally Private Heavy Hitters

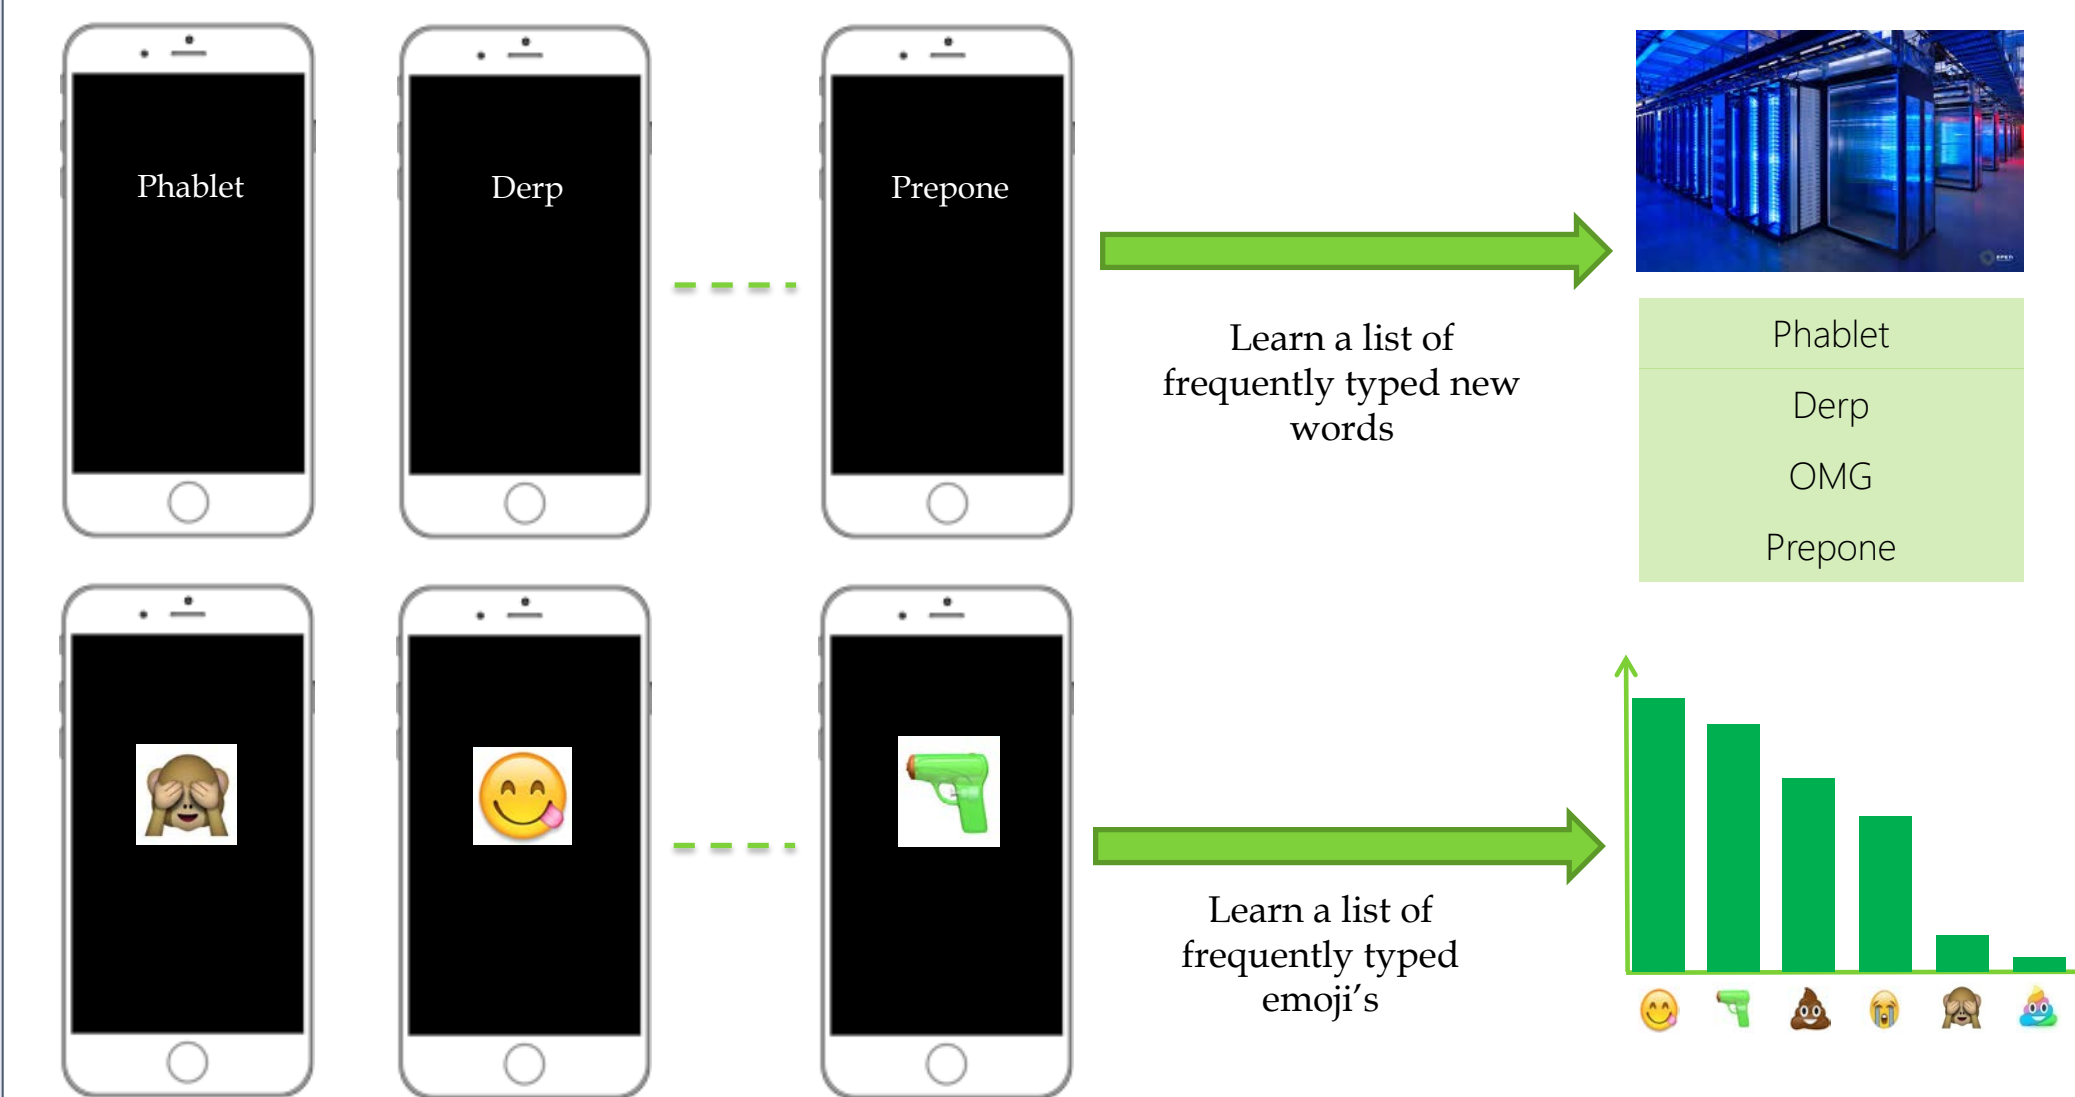
Raef Bassily<sup>1</sup> Kobbi Nissim<sup>2</sup> Uri Stemmer<sup>3</sup> Abhradeep Thakurta<sup>4</sup>

<sup>1</sup>Ohio State University <sup>2</sup>Gerogetown University <sup>3</sup>Harvard University <sup>4</sup>University of California Santa Cruz



## Learning from private data


Example: Learn new words and emoji's



### Problem definition

- Consider a set of data samples distributed across devices
  - Each data sample corresponds to one device
- Want: Learn *frequency statistics* and *heavy hitters* while preserving *local differential privacy*
  - Heavy hitters:** All elements in the domain with frequency higher than a fixed threshold

Local differential privacy [Warner 65, EGS 03, DMNS 06, KLNRS 08]



Intuition: Every possible input  $x$  is plausible given transcript  $\mathcal{A}(x)$ . Need not trust a centralized data curator.

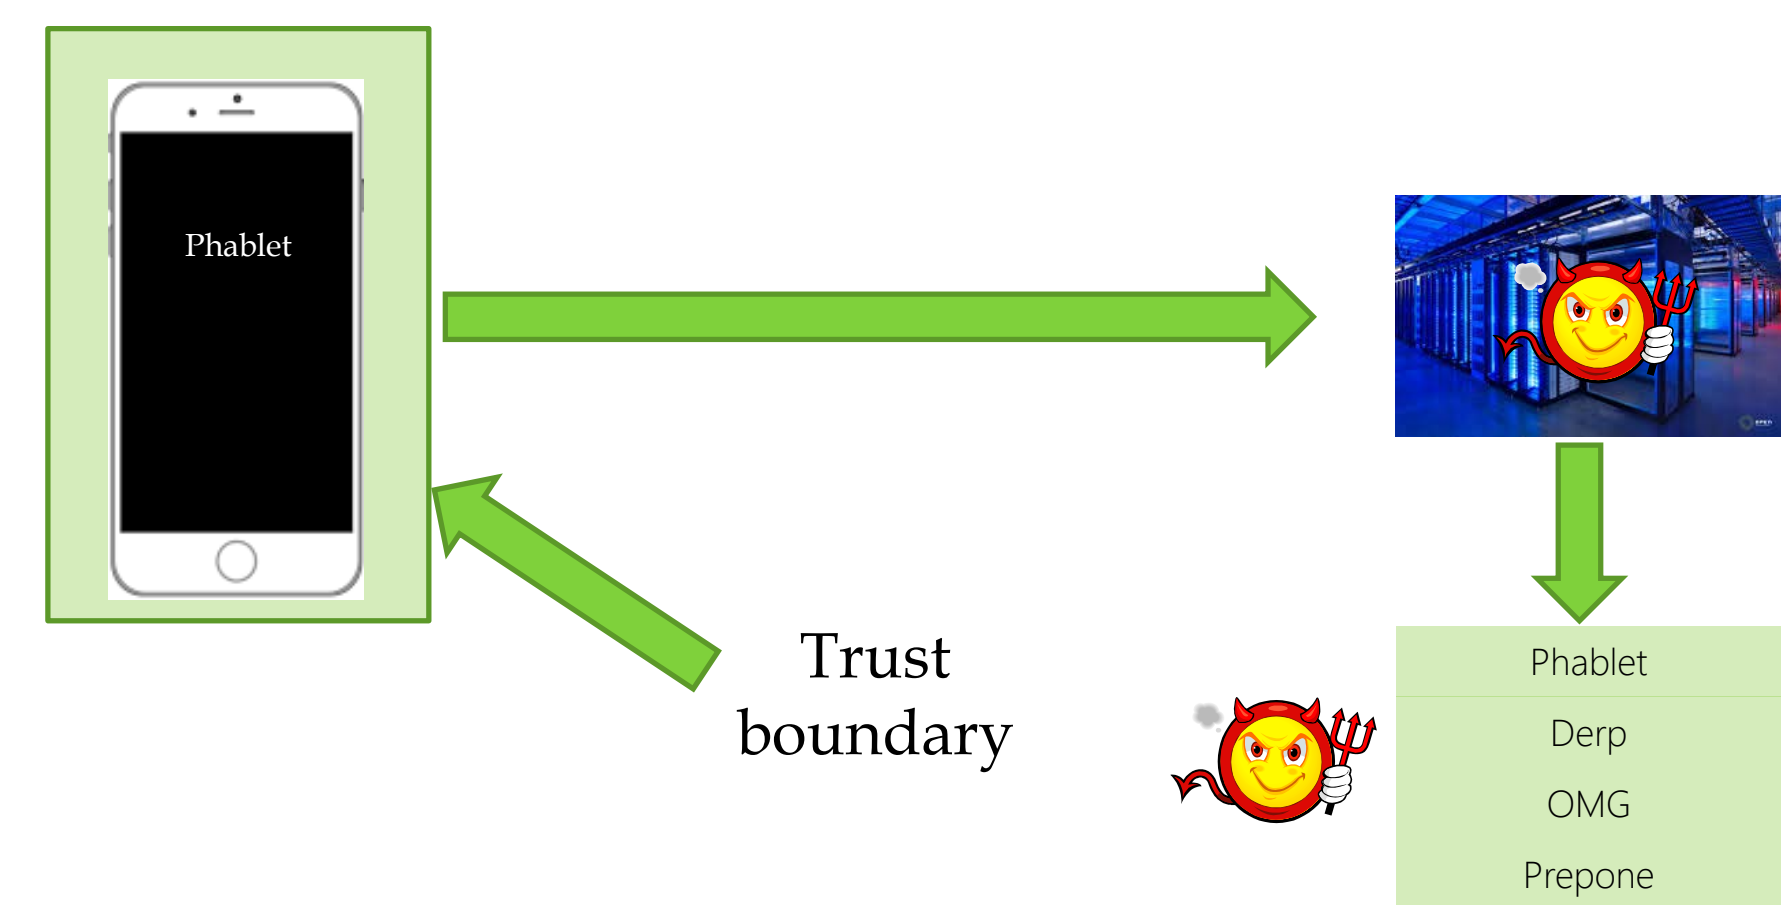
A randomized algorithm  $\mathcal{A}$  is  $\epsilon$ -locally differentially private if

- for all possible data samples  $x, x'$  from the domain
- for all (measurable) sets of answers  $E$

$$\Pr[\mathcal{A}(x) \in E] \leq e^\epsilon \Pr[\mathcal{A}(x') \in E]$$

## Why is privacy a concern?

- Adversary can *observe the output*, or even worse, *compromise the server*
- Companies use similar technology for analyzing *very sensitive data*
  - Google uses it for collecting malicious URL visit
  - Apple uses it for collecting new words people type, health related information
  - LinkedIn uses it for collecting salary information



## Our contributions

Two new algorithms with optimal error, with optimal time, space and communication complexity

- TreeHist:** Based on classic count sketch [CCFC02]
  - Nearly optimal error guarantee
  - Easily implementable in practice
- Bitstogram:** Based on error correcting codes
  - Optimal error guarantee

Implementation of our TreeHist protocol

- Scales to 10 million data samples, with domain  $\approx 10^8$
- Performs better than state-of-the-art open source tools

Our results at a glance

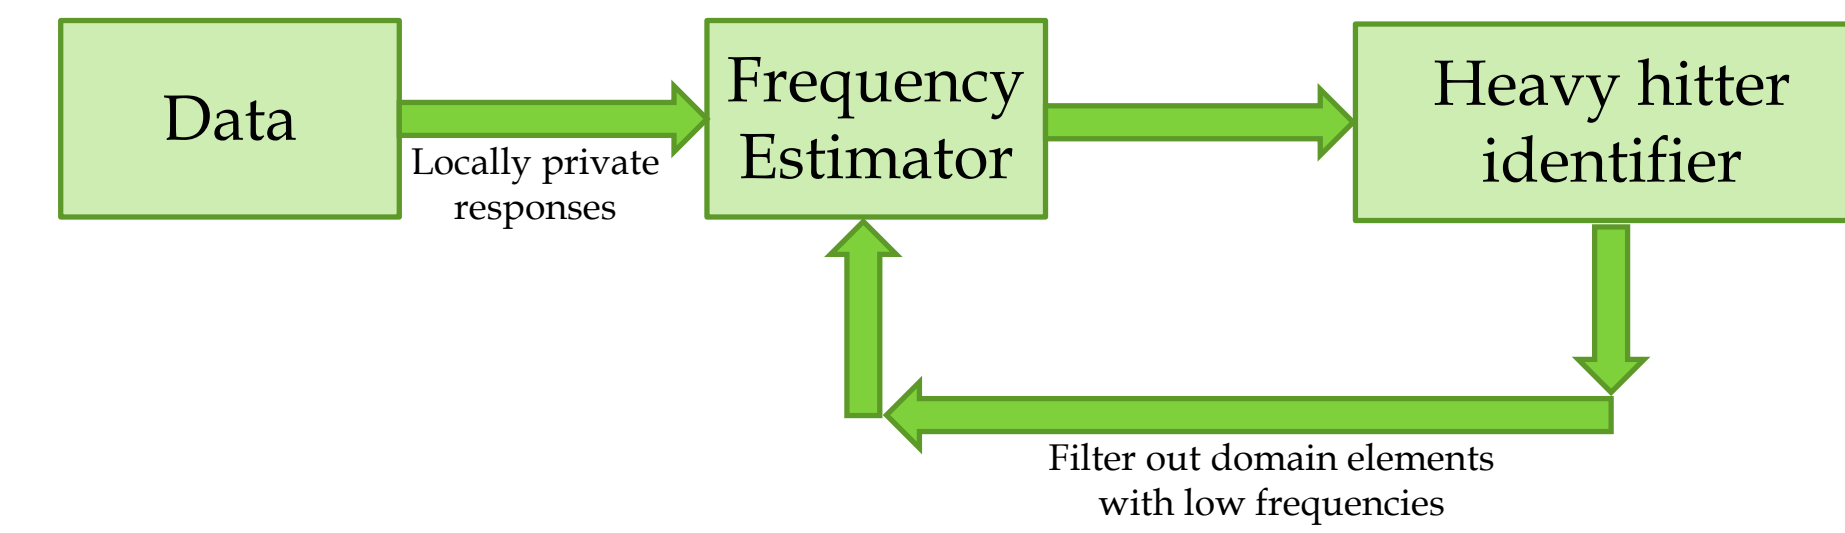
Performance metric	TreeHist (this work)	Bitstogram (this work)	Bassily and Smith'15
Server time	$\tilde{O}(n)$	$\tilde{O}(n)$	$\tilde{O}(n^{5/2})$
User time	$\tilde{O}(1)$	$\tilde{O}(1)$	$\tilde{O}(n^{3/2})$
Server memory	$\tilde{O}(\sqrt{n})$	$\tilde{O}(\sqrt{n})$	$\tilde{O}(n^2)$
User memory	$\tilde{O}(1)$	$\tilde{O}(1)$	$\tilde{O}(n^{3/2})$
Communication/user	$\tilde{O}(1)$	$\tilde{O}(1)$	$\tilde{O}(1)$
Error	$o(\sqrt{n \log n \log d})$	$o(\sqrt{n \log d})$	$o(\sqrt{n \log d})$

$n$ : Number of data samples  
 $d$ : Domain size  
 Error  $\eta$ : Can find all heavy hitters with frequency  $\geq 3\eta$ , and none with frequency  $< \eta$   
 (Table hides terms in  $\epsilon$ )

## Prior work

- [Bassily and Smith 2015]
  - Algorithm based on dimensionality reduction via the Johnson Lindenstrauss transformation
  - Optimal error guarantee, non-optimal in terms of computation, storage and communication
- [Thakurta et al., 2016]
  - Optimal frequency estimator via count-median-sketch
  - Heavy hitters algorithm is a heuristic, with no formal guarantees
- [Erlingsson et al., 2013, Fanti et al., 2016]
  - Empirical results only, with no formal guarantees
- [Hsu et al., 2014]
  - Algorithm based on group testing

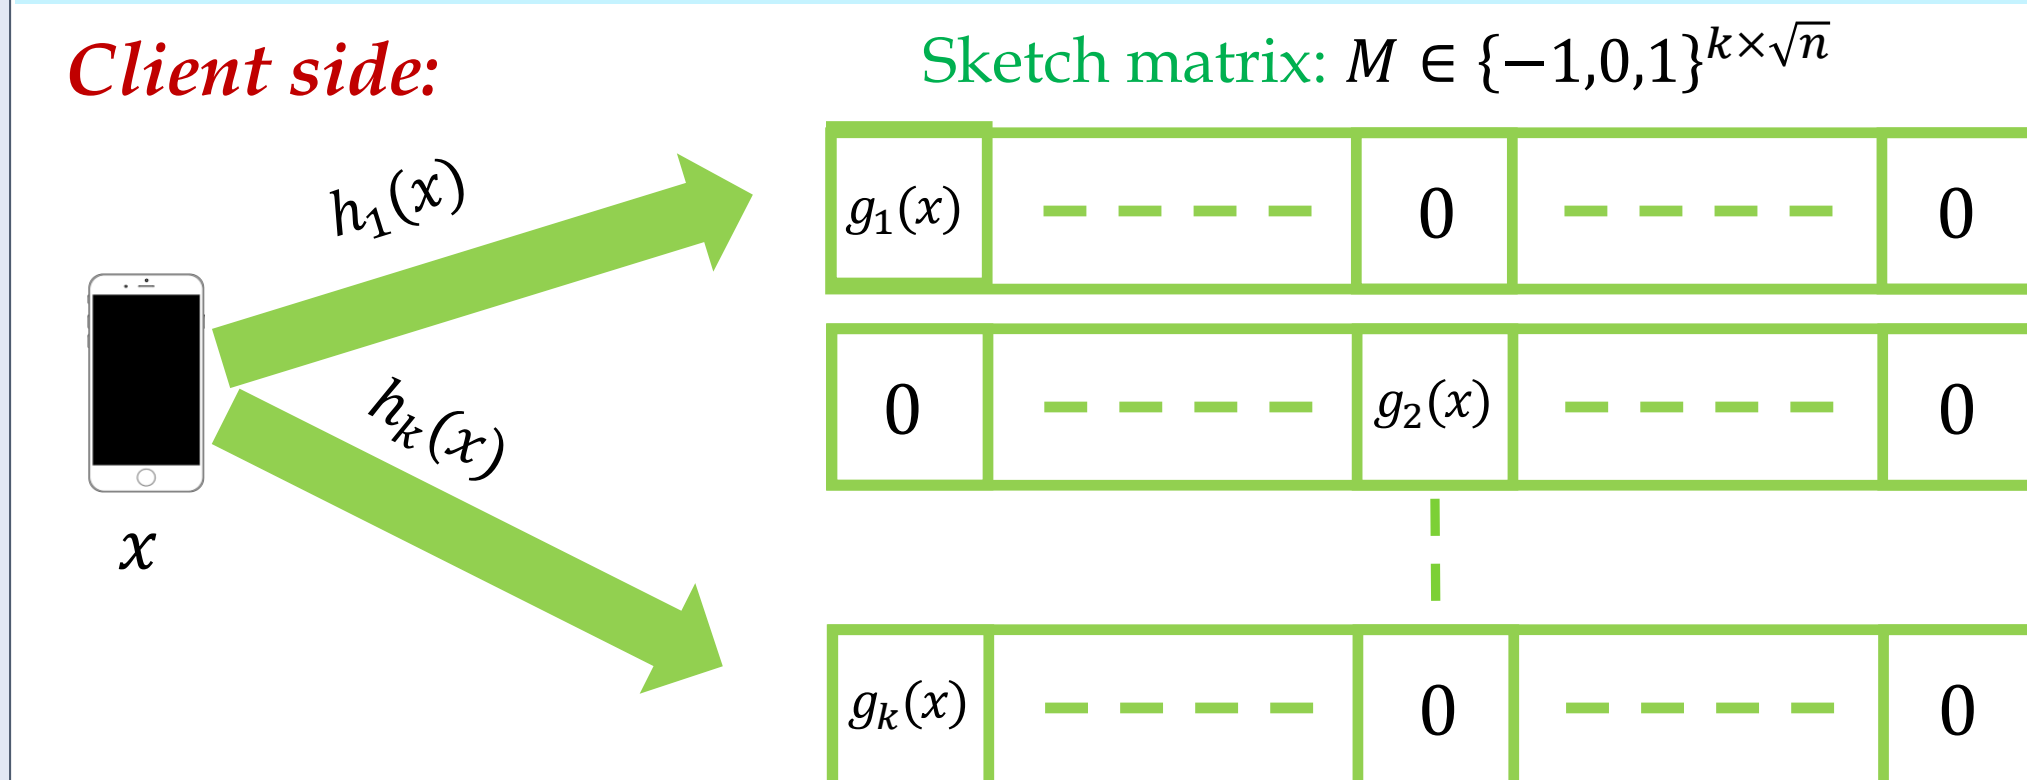
## Algorithm overview



### Locally private frequency estimator

**Problem:** Design locally differentially private estimator with  
 Accuracy:  $O(\sqrt{n \log d})$  Query time / estimate:  $O(\log d)$   
 Communication: 1 bit/user Client computation: Constant  
 Server computation:  $O(\sqrt{n \log d})$

Non-private count-sketch [CCFC02]



Hash functions: Pairwise independent s.t.  $\forall i \in [k], h_i: [d] \rightarrow [\sqrt{n}]$   
 Pairwise independent s.t.  $\forall i \in [k], g_i: [d] \rightarrow \{\pm 1\}$



Frequency estimate of  $x$ : Median $\{g_i(x)S[i, h_i(x)]\}$

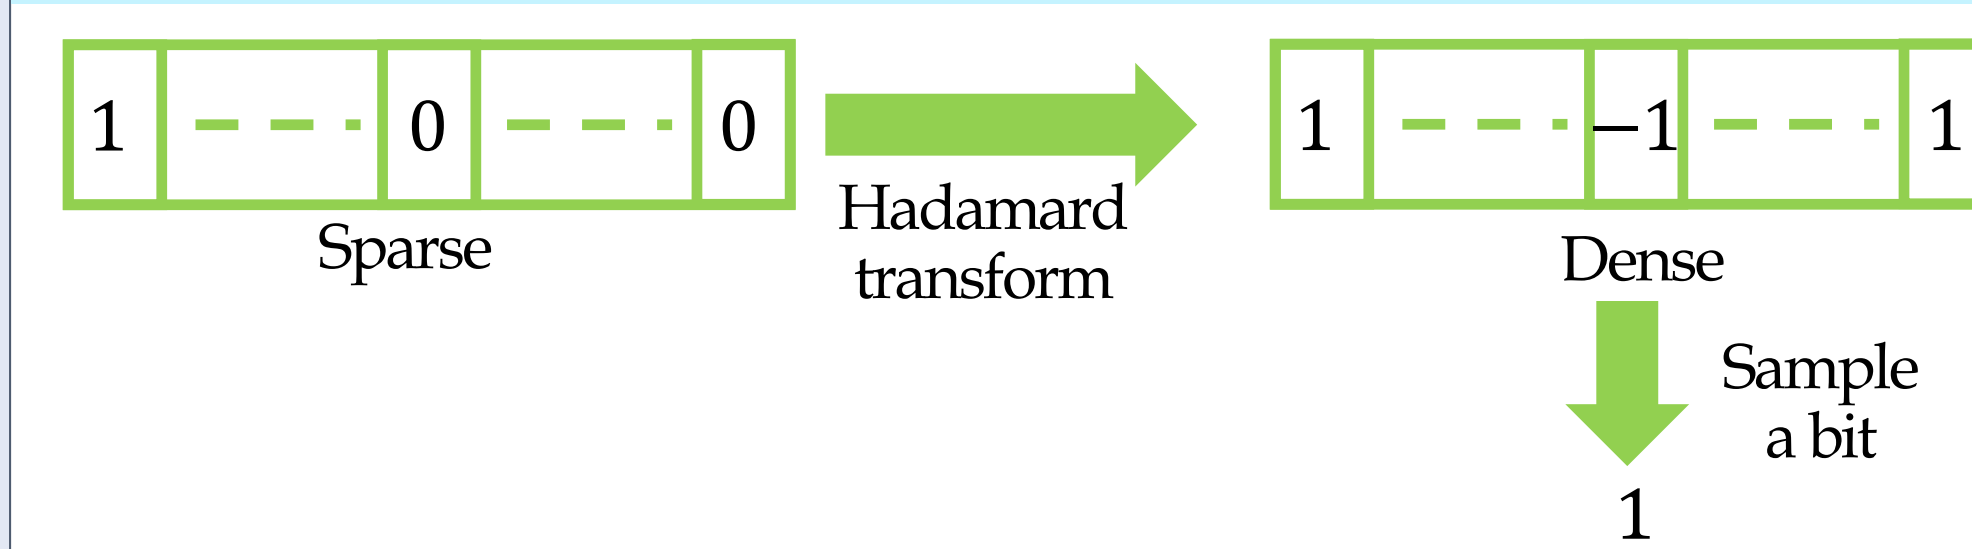
### Differentially private count sketch

- Client side:
- Sample a row  $M_i$  from the sketch matrix  $M$
  - Replace each zero of  $M_i$  with  $\pm 1$  sampled w.p. =  $1/2$
  - Flip each bit of  $M_i$  with probability  $\frac{e^\epsilon}{1+e^\epsilon}$ , and send  $(i, M_i)$

Server side: For client report  $(i, v)$ , add  $(\log d \times \frac{e^\epsilon + 1}{e^\epsilon - 1}) \cdot v$  to the  $i$ -th row of the server side sketch  $S$

**Theorem:** The algorithm is  $\epsilon$  locally differentially private, and satisfies all the desired utility guarantees, *except communication*

### Reducing communication via Hadamard transform




**Theorem:** Hadamard transform and sampling does not affect privacy or utility, but *reduces communication to  $\tilde{O}(1)$  bit*

## (Simplified) TreeHist protocol

Observation: Any heavy hitter can be represented using  $\log d$  bits

Idea: Construct each heavy hitter bit by bit via its prefix



Frequent prefixes of length two

Observation: If a domain element has frequency  $\geq \sqrt{n}$ , so does all its prefixes

Observation: Roughly  $\sqrt{n}$  elements per level of the tree

Algorithm: Each client samples a random prefix, and sends it via the locally differentially private count sketch protocol

The server computes the frequent heavy hitters via the prefix tree

## (Simplified) Bitstogram protocol

- Denote  $S = (x_1, \dots, x_n)$  where  $x_i \in [d]$  is the data of user  $i$
- Assume we have frequency oracle  $\mathcal{O}$  with error  $\tau$
- Public randomness: Hash function  $h: X \rightarrow [T]$   
 Intuition: if  $h$  isolates heavy-hitters then suffices to query  $\mathcal{O}$  on hash range. But how?
- Fix a "heavy-hitter"  $x^*$  appearing  $2\tau$  times in the data  $S$ . Denote  $t^* = h(x^*)$ .  
 For simplicity: assume that only  $x^*$  maps into  $t^*$  (from  $S$ )  
 Observe: every user can compute her row locally

$x^*$  is "heavy", hence  $(t^*, x^*[\ell])$  appears  $> 2\tau$  in  $S_\ell$  for every  $\ell$   
 No collisions, hence  $(t^*, 1 - x^*[\ell])$  appears 0 times in  $S_\ell$   
 $\Rightarrow$  Can identify every bit  $\ell$  of  $x^*$  by querying  $\mathcal{O}(S_\ell)$  on  $(t^*, 0)$  and  $(t^*, 1)$

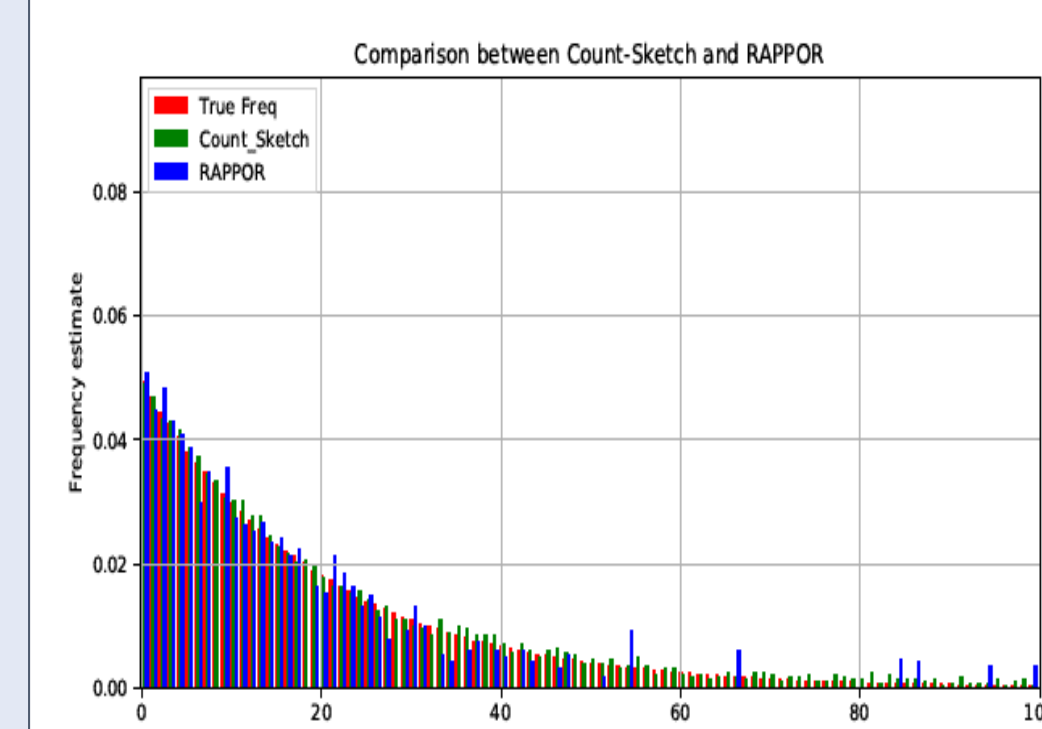
**The Protocol:** For every  $t \in [T]$  construct  $\hat{x}^{(t)}$  as follows:  
 $\forall \ell \in [\log |X|]$  query  $\mathcal{O}(S_\ell)$  on  $(t, 0)$  and  $(t, 1)$  and set  $\hat{x}^{(t)}[\ell] \leftarrow \text{argmax}$

- By purple,  $\hat{x}^{(t^*)} = x^*$  is identified

## Empirical evaluation with TreeHist

Comparison to RAPPOR project:  $\epsilon = 1.09, n = 1 \text{ mi}, \text{domain} = 100$

Heavy hitters for NLTK data set:  $\epsilon = 2, n = 10 \text{ mi}, \# \text{ of unique words} = 25991, \text{wordlength} = 6$   
 Threshold =  $15\sqrt{n}$



True positives: 22	True negatives: $\approx 3 \times 10^8$
False positives: 60	False negatives: 3