

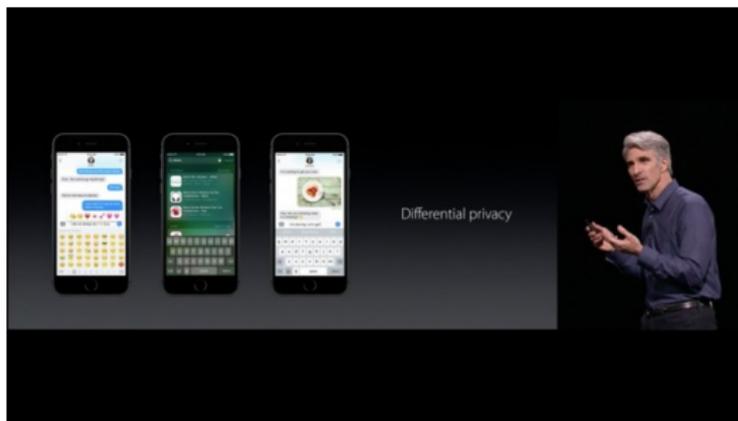
Formal verification of differentially private computations

Gilles Barthe
IMDEA Software Institute, Madrid, Spain

July 13, 2016

Apple WorldWide Developer Conference

San Francisco, June 13-17, 2016



Data analysis

For companies:

- ▶ Know users
- ▶ Provide better services
- ▶ Reduce fraud



For health organizations:

- ▶ Establish genetic correlations
- ▶ Monitor epidemia
- ▶ Take actions



For users: *your data, but not your data*

- ▶ Accurate computations
- ▶ Individual privacy



Data anonymization

Name	Gender	Birth data	Zip code	Status
Ann	F	12/01/1957	28000	Y
Bob	M	25/06/1970	28006	X
Eve	F	13/05/1953	28001	X
Tom	M	06/11/1966	28001	Y

Data anonymization

Name	Gender	Birth data	Zip code	Status
bafkb	F	12/01/1957	28000	Y
kaau	M	25/06/1970	28006	X
jmbag	F	13/05/1953	28001	X
cjhuf	M	06/11/1966	28001	Y

Can infer Tom's status from knowing

- ▶ Tom is in the database
- ▶ his gender, birth date, zip code

Data anonymization

Name	Gender	Birth data	Zip code	Status
bafkb	F	12/01/1957	28000	Y
kaau	M	25/06/1970	28006	X
jmbag	F	13/05/1953	28001	X
cjhuf	M	06/11/1966	28001	Y

Can infer Tom's status from knowing

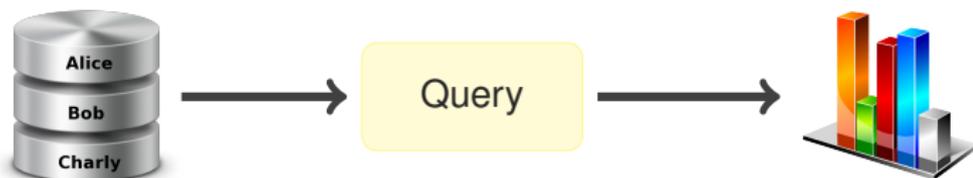
- ▶ Tom is in the database
- ▶ his gender, birth date, zip code

A real problem

- ▶ Uniquely identify $\geq 85\%$ of individuals
- ▶ Adding noise does not help

Differential privacy

(Dwork, McSherry, Nissim, and Smith, 2006)



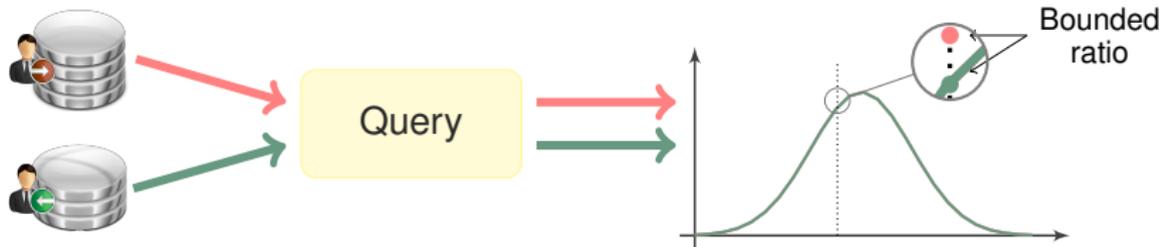
Differential privacy

(Dwork, McSherry, Nissim, and Smith, 2006)



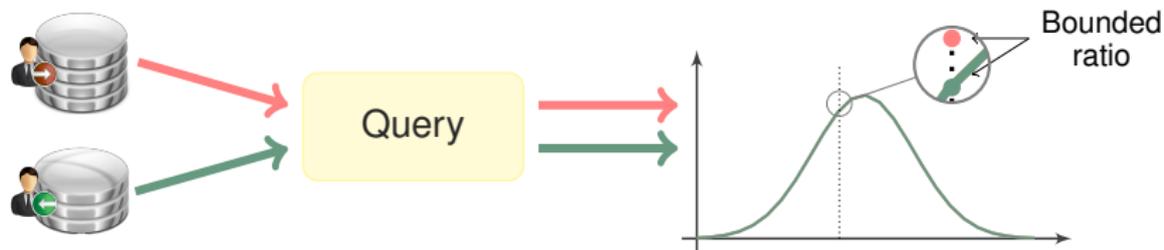
Differential privacy

(Dwork, McSherry, Nissim, and Smith, 2006)



Differential privacy

(Dwork, McSherry, Nissim, and Smith, 2006)

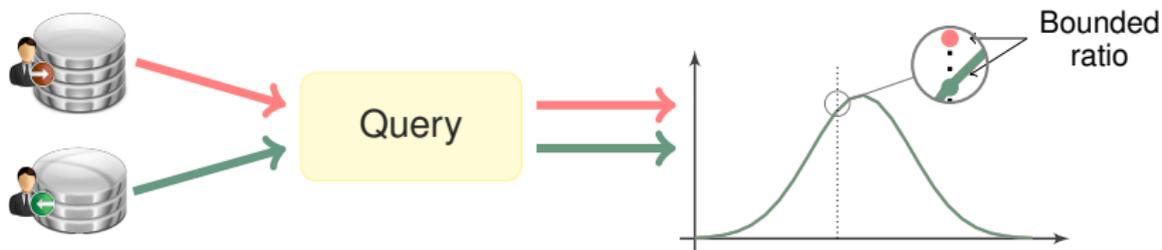


A randomized algorithm \mathcal{K} is (ϵ, δ) -differentially private w.r.t. Φ iff for all databases D_1 and D_2 s.t. $\Phi(D_1, D_2)$

$$\forall S. \Pr[\mathcal{K}(D_1) \in S] \leq \exp(\epsilon) \cdot \Pr[\mathcal{K}(D_2) \in S] + \delta$$

Differential privacy

(Dwork, McSherry, Nissim, and Smith, 2006)



A randomized algorithm \mathcal{K} is (ϵ, δ) -differentially private w.r.t. Φ iff for all databases D_1 and D_2 s.t. $\Phi(D_1, D_2)$

$$\forall S. \Pr[\mathcal{K}(D_1) \in S] \leq \exp(\epsilon) \cdot \Pr[\mathcal{K}(D_2) \in S] + \delta$$

Special case

If $\epsilon \approx 0$ and $\delta = 0$, then it suffices that for all databases D_1 and D_2 s.t. $\Phi(D_1, D_2)$

$$\forall x. \Pr[\mathcal{K}(D_1) = x] \approx (1 + \epsilon) \cdot \Pr[\mathcal{K}(D_2) = x]$$

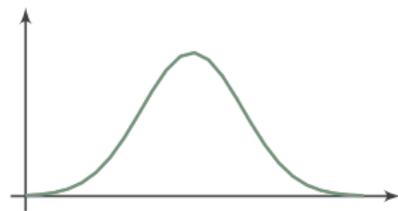
Advantages of differential privacy

- ▶ Mathematically rigorous
- ▶ Many algorithms have a private and accurate realization
- ▶ Mechanisms for achieving privacy via output perturbation
- ▶ Composition theorems for private algorithms

Differential privacy via output perturbation

Let f be k -sensitive w.r.t. Φ :

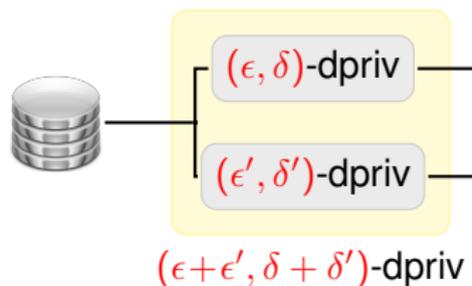
$$\Phi(a, a') \implies |f(a) - f(a')| \leq k$$



Then $a \mapsto \text{Lap}_\epsilon(f(a))$ is $(k \cdot \epsilon, 0)$ -differentially private w.r.t. Φ

Differential privacy by sequential composition

- ▶ If \mathcal{K} is (ϵ, δ) -differentially private, and
- ▶ $\lambda a. \mathcal{K}'(a, b)$ is (ϵ', δ') -differentially private for every $b \in B$,
- ▶ then $\lambda a. \mathcal{K}'(a, \mathcal{K}(a))$ is $(\epsilon + \epsilon', \delta + \delta')$ -differentially private



Application: noisy sums

```
function NOISYSUM1(a)  
s = 0; i = 0;  
while i < length(a) do  
    s = s + a[i];  
    i = i + 1;  
end;  
s = Lapε(s);  
return s
```

NOISYSUM₁ is $(b \cdot \epsilon, 0)$ -differentially private

Application: noisy sums

```
function NOISYSUM2( $a$ )  
   $s = 0; i = 0;$   
  while  $i < \text{length}(a)$  do  
     $\tilde{a} = \text{Lap}_\epsilon(a[i]);$   
     $s = s + \tilde{a};$   
     $i = i + 1;$   
  end;  
  return  $s$ 
```

NOISYSUM₂ is $(n \cdot b \cdot \epsilon, 0)$ -differentially private

Differential privacy beyond sequential composition

There is much more to differential privacy

- ▶ Exponential mechanism
- ▶ Optimal composition
- ▶ Adaptive adversaries
- ▶ Accuracy-dependent privacy
- ▶ Also, many variants of differential privacy

Issues

- ▶ Proofs are intricate and may be wrong
- ▶ Proofs, when correct, are messy
- ▶ Hard to predict when altering an algorithm breaks privacy

The Sparse Vector Technique

```
SparseVectorbt(a, b, M, N, d) :=  
  i = 0; l = [];  
  u  $\stackrel{\$}{\leftarrow}$  Lapε(0); A = a - u; B = b + u;  
  while i < N do  
    q = A(l);  
    S  $\stackrel{\$}{\leftarrow}$  Lapε(evalQ(q, d));  
    if (A ≤ S ≤ B ∧ |l| < M) then l = i :: l;  
    i = i + 1;  
  return l
```

IF all queries are **1**-sensitive,
THEN algorithm achieves $(\sqrt{M}\epsilon, \delta')$ -differential privacy
EVEN IF $M \ll N$

3.1 Privacy Proof for Algorithm 1

We now prove the privacy of Algorithm 1. We break the proof down into two steps, to make the proof easier to understand, and more importantly, to point out what confusions likely caused the different non-private variants of SVT to be proposed. In the first step, we analyze the situation where the output is \perp^L , a length- L vector (\perp, \dots, \perp) , indicating that all l queries are found to be below the threshold.

LEMMA 1. Let A be Algorithm 1. For any neighboring datasets D and D' , and any integer l , we have

$$\Pr[A(D) = \perp^L] \leq e^\epsilon \Pr[A(D') = \perp^L].$$

PROOF. We have

$$\Pr[A(D) = \perp^L] = \int_{z \in \mathbb{R}} f_L(D, z, L) dz,$$

$$\text{where } f_L(D, z, L) = \Pr[\rho = z] \prod_{i=1}^L \Pr[\eta_i(D) + v_i < T_i + z], \quad (1)$$

and $L = \{1, 2, \dots, l\}$.

The probability of outputting \perp^L over D is the summation (or integral) of terms $f_L(D, z, L)$, each of which is the product of $\Pr[\rho = z]$, the probability that the threshold noise equals z , and $\prod_{i=1}^L \Pr[\eta_i(D) + v_i < T_i + z]$, the conditional probability that \perp^L is the output on D given that the threshold noise is z . (Note that given D , T , the queries, and ρ , whether one query results in \perp or not depends completely on the noise v_i and is independent from whether any other query results in \perp .) If we can prove

$$f_L(D, z, L) \leq e^\epsilon f_L(D', z + \Delta, L), \quad (2)$$

then we have

$$\begin{aligned} \Pr[A(D) = \perp^L] &= \int_{z \in \mathbb{R}} f_L(D, z, L) dz \\ &\leq \int_{z \in \mathbb{R}} e^\epsilon f_L(D', z + \Delta, L) dz \quad \text{from (2)} \\ &\leq e^\epsilon \int_{z \in \mathbb{R}} f_L(D', z', L) dz' \quad \text{let } z' = z + \Delta \\ &= e^\epsilon \Pr[A(D') = \perp^L]. \end{aligned}$$

This proves the lemma. It remains to prove Eq (2). For any query q_i , because $|q_i(D) - q_i(D')| \leq \Delta$ and thus $-q_i(D) \leq \Delta - q_i(D')$, we have

$$\begin{aligned} \Pr[\eta_i(D) + v_i < T_i + z] &= \Pr[v_i < T_i - q_i(D) + z] \\ &\leq \Pr[v_i < T_i + \Delta - q_i(D') + z] \\ &= \Pr[q_i(D') + v_i < T_i + (z + \Delta)] \quad (3) \end{aligned}$$

With (3), we prove (2) as follows:

$$\begin{aligned} f_L(D, z, L) &= \Pr[\rho = z] \prod_{i \in L} \Pr[\eta_i(D) + v_i < T_i + z] \\ &\leq e^\epsilon \Pr[\rho = z + \Delta] \prod_{i \in L} \Pr[q_i(D') + v_i < T_i + (z + \Delta)] \\ &= e^\epsilon f_L(D', z + \Delta, L). \end{aligned}$$

□

That is, by using a noisy threshold, we are able to bound the probability ratio for all the negative query answers (i.e., \perp^L) by e^ϵ , no matter how many negative answers there are.

We can obtain a similar result for positive query answers in the same way.

$$\text{Let } f_T(D, z, L) = \Pr[\rho = z] \prod_{i \in L} \Pr[\eta_i(D) + v_i \geq T_i + z].$$

We have $f_T(D, z, L) \leq e^\epsilon f_T(D', z - \Delta, L)$, and thus

$$\Pr[A(D) = T^L] \leq e^\epsilon \Pr[A(D') = T^L].$$

This likely contributes to the misunderstandings behind Algorithms 5 and 6, which treat positive and negative answers exactly the same way. The problem is that while one is free to choose to bound positive or negative side, one cannot bound both.

We also observe that the proof of Lemma 1 will go through if no noise is added to the query answers, i.e., $v_i = 0$, because Eq (3) holds even when $v_i = 0$. It is likely because of this observation that Algorithm 5 adds no noise to query answers. However, when considering outcomes that include both positive answers (T^L) and negative answers (\perp^L), one has to add noises to the query answers, as we show below.

THEOREM 2. Algorithm 1 is ϵ -DP.

PROOF. Consider any output vector $a \in \{\perp, T\}^L$. Let $\alpha = (\alpha_1, \dots, \alpha_l)$, $\eta^{\perp} = \{i : a_i = \perp\}$, and $\eta^T = \{i : a_i = T\}$. Clearly, $|\eta^{\perp}| \leq l$. We have

$$\begin{aligned} \Pr[A(D) = a] &= \int_{z \in \mathbb{R}} g(D, z) dz, \text{ where} \\ g(D, z) &= \Pr[\rho = z] \prod_{i \in \eta^{\perp}} \Pr[\eta_i(D) + v_i < T_i + z] \prod_{i \in \eta^T} \Pr[\eta_i(D) + v_i \geq T_i + z]. \end{aligned}$$

We want to show that $g(D, z) \leq e^\epsilon g(D', z + \Delta)$. This suffices to prove that $\Pr[A(D) = a] \leq e^\epsilon \Pr[A(D') = a]$. Note that $g(D, z)$ can be written as:

$$g(D, z) = f_L(D, z, \eta^{\perp}) \prod_{i \in \eta^T} \Pr[\eta_i(D) + v_i \geq T_i + z].$$

Following the proof of Lemma 1, we can show that $f_L(D, z, \eta^{\perp}) \leq e^\epsilon f_T(D', z + \Delta, \eta^{\perp})$, and it remains to show

$$\prod_{i \in \eta^T} \Pr[\eta_i(D) + v_i \geq T_i + z] \leq e^\epsilon \prod_{i \in \eta^T} \Pr[\eta_i(D') + v_i \geq T_i + z + \Delta]. \quad (4)$$

Because $v_i = \text{Lap}(\frac{\epsilon \Delta}{2})$ and $|\eta_i(D) - \eta_i(D')| \leq \Delta$, we have

$$\begin{aligned} \Pr[\eta_i(D) + v_i \geq T_i + z] &= \Pr[v_i \geq T_i + z - \eta_i(D)] \\ &\leq \Pr[v_i \geq T_i + z - \Delta - \eta_i(D')] \quad (5) \\ &\leq e^\epsilon \Pr[v_i \geq T_i + z - \Delta - \eta_i(D') + 2\Delta] \quad (6) \\ &= e^\epsilon \Pr[\eta_i(D') + v_i \geq T_i + z + \Delta]. \end{aligned}$$

Eq (5) is because $-q_i(D) \geq \Delta - q_i(D')$, and Eq (6) is from the Laplace distribution's property. This proves Eq (4). □

The basic idea of the proof is that when computing $g(D, z)$ with $g(D', z + \Delta)$, we can bound the probability ratio for all outputs of \perp to no more than e^ϵ by using a noisy threshold, no matter how many such outputs there are. To bound the ratio for the T outputs to no more than e^ϵ , we need to add sufficient Laplacian noises, which should scale with ϵ , the number of positive outputs.

Now we turn to Algorithm 3-6 to clarify what are wrong with their privacy proofs and to give their DP properties.

3.1 Privacy Proof for Algorithm 1

We now prove the privacy of Algorithm 1. We break the proof down into two steps, to make the proof easier to understand, and more importantly, to point out what confusions likely caused the different non-private variants of SVT to be proposed. In the first step, we analyze the situation where the output is ± 1 , a length- L vector (x_1, \dots, x_L) , indicating that all L queries are tested to be below the threshold.

LEMMA 1. Let A be Algorithm 1. For any neighboring datasets D and D' , and any integer L , we have

$$\Pr[A(D) = \pm 1] \leq e^{\epsilon} \Pr[A(D') = \pm 1].$$

PROOF. We have

$$\Pr[A(D) = \pm 1] = \int_{\pm 1} f_L(D, z, L) dz,$$

$$\text{where } f_L(D, z, L) = \Pr[\rho = \pm 1] \prod_{i=1}^L \Pr[\rho_i(D) + v_i < T_i + z], \quad (1)$$

and $L = \{1, 2, \dots, L\}$.

The probability of outputting ± 1 over D is the summation (or integral) of terms $f_L(D, z, L)$, each of which is the product of $\Pr[\rho = \pm 1]$, the probability that the threshold noise equals z , and $\prod_{i=1}^L \Pr[\rho_i(D) + v_i < T_i + z]$, the conditional probability that ± 1 is the output on D given that the threshold noise is z . (Note that given D , T , the queries, and ρ , whether one query results in ± 1 or not depends completely on the noise v_i and is independent from whether any other query results in ± 1 .) If we can prove

$$f_L(D, z, L) \leq e^{\epsilon} f_L(D', z, \Delta, L), \quad (2)$$

then we have

$$\begin{aligned} \Pr[A(D) = \pm 1] &= \int_{\pm 1} f_L(D, z, L) dz \\ &\leq \int_{\pm 1} e^{\epsilon} f_L(D', z, \Delta, L) dz \quad \text{from (2)} \\ &\leq e^{\epsilon} \int_{\pm 1} f_L(D', z, L) dz \quad \text{let } z' = z + \Delta \\ &= e^{\epsilon} \Pr[A(D') = \pm 1]. \end{aligned}$$

This proves the lemma. It remains to prove Eq (2). For any query q_i , we have $|q_i(D) - q_i(D')| \leq \Delta$ and thus $-q_i(D) \leq -\Delta - q_i(D')$.

$$\begin{aligned} \Pr[\rho_i(D) + v_i < T_i + z] &= \Pr[\rho_i < T_i - q_i(D) + z] \\ &\leq \Pr[v_i < T_i - \Delta - q_i(D') + z] \\ &= \Pr[q_i(D') + v_i < T_i + (z + \Delta)] \quad (3) \end{aligned}$$

With (3), we prove (2) as follows:

$$\begin{aligned} f_L(D, z, L) &= \Pr[\rho = \pm 1] \prod_{i \in L} \Pr[\rho_i(D) + v_i < T_i + z] \\ &\leq e^{\epsilon} \Pr[\rho = \pm 1] \prod_{i \in L} \Pr[\rho_i(D') + v_i < T_i + (z + \Delta)] \\ &= e^{\epsilon} f_L(D', z + \Delta, L). \end{aligned}$$

□

That is, by using a noisy threshold, we are able to bound the privacy for all of the negative query answers (i.e., ± 1) by e^{ϵ} , no matter how many negative answers there are.

We can obtain a similar result for positive query answers in the same way.

$$\text{Let } f_L(D, z, L) = \Pr[\rho = \pm 1] \prod_{i \in L} \Pr[\rho_i(D) + v_i > T_i + z].$$

$$\begin{aligned} \text{We have } f_L(D, z, L) &\leq e^{\epsilon} f_L(D', z - \Delta, L), \text{ and thus} \\ \Pr[A(D) = \mp 1] &\leq e^{\epsilon} \Pr[A(D') = \mp 1]. \end{aligned}$$

This likely contributes to the misunderstandings behind Algorithms 5 and 6, which treat positive and negative answers exactly the same way. The problem is that while one is free to choose to bound positive or negative side, one cannot bound both.

We also observe that the proof of Lemma 1 will go through if no noise is added to the query answers, i.e., $v_i = 0$. Because Eq (3) holds even when $v_i = 0$, it is likely because of this observation that Algorithm 5 adds no noise to query answers. However, when considering outcomes that include both positive answers (T 's) and negative answers (L 's), one has to add noises to the query answers, as we show below.

THEOREM 2. Algorithm 1 is ϵ -DP.

PROOF. Consider any output vector $a \in \{\pm 1, \top\}^L$. Let $a = (a_1, \dots, a_L)$, $\pi^+ = \{i : a_i = \top\}$, and $\pi^- = \{i : a_i = \pm 1\}$. Clearly, $|\pi^+| \leq L$. We have

$$\begin{aligned} \Pr[A(D) = a] &= \int_{a} f_L(D, z, L) dz, \text{ where} \\ \rho_i(D, z) &= \Pr[\rho = \pm 1] \prod_{i \in \pi^+} \Pr[\rho_i(D) + v_i < T_i + z] \prod_{i \in \pi^-} \Pr[\rho_i(D) + v_i > T_i + z]. \end{aligned}$$

We want to show that $\rho_i(D, z) \leq e^{\epsilon} \rho_i(D', z + \Delta)$. This suffices to prove that $\Pr[A(D) = a] \leq e^{\epsilon} \Pr[A(D') = a]$. Note that $A(D, z)$ can be written

$$\rho_i(D, z) = f_L(D, z, L) \prod_{i \in \pi^+} \Pr[\rho_i(D) + v_i \geq T_i + z].$$

Following the proof of Lemma 1, we can show that $f_L(D, z, L) \leq e^{\epsilon} f_L(D', z + \Delta, L)$, and it remains to show

$$\prod_{i \in \pi^+} \Pr[\rho_i(D) + v_i \geq T_i + z] \leq e^{\epsilon} \prod_{i \in \pi^+} \Pr[\rho_i(D') + v_i \geq T_i + z + \Delta]. \quad (4)$$

Because $v_i = \text{Lap}(\frac{\Delta}{2\epsilon})$ and $|\rho_i(D) - \rho_i(D')| \leq \Delta$, we have

$$\begin{aligned} \Pr[\rho_i(D) + v_i \geq T_i + z] &= \Pr[\rho_i \geq T_i + z - q_i(D)] \\ &\leq \Pr[v_i \geq T_i + z - \Delta - q_i(D')] \quad (5) \\ &\leq e^{\epsilon} \Pr[v_i \geq T_i + z - \Delta - q_i(D') + \Delta] \quad (6) \\ &= e^{\epsilon} \Pr[v_i \geq T_i + z + \Delta]. \end{aligned}$$

Eq (5) is because $-q_i(D) \geq -\Delta - q_i(D')$, and Eq (6) is from the Laplace distribution's property. This proves Eq (4). □

The basic idea of the proof is that when comparing $\rho_i(D, z)$ with $\rho_i(D', z + \Delta)$, we can bound the probability ratio for all outputs of ± 1 to no more than e^{ϵ} by using a noisy threshold, no matter how many such outputs there are. To bound the ratio for the T outputs to no more than e^{ϵ} , we need to add sufficient Laplacian noises, which should scale with ϵ , the number of positive outputs.

Now we turn to Algorithm 3-6 to clarify what are wrong with their privacy proofs and to give their DP properties.

Figure 1: A Selection of SVT Variants

Input/Output shared by all SVT Algorithms

Input: A private database D , a stream of queries $Q = q_1, q_2, \dots$ each with sensitivity no more than Δ , either a sequence of thresholds $T = T_1, T_2, \dots$ or a single threshold T (use footnote ¹), and ϵ , the maximum number of queries to be answered with T .
Output: A stream of answers a_1, a_2, \dots , where each $a_i \in \{T, \pm 1\}$ and \mathbb{R} denotes the set of all real numbers.

Algorithm 1 An instantiation of the SVT proposed in this paper

```
Input:  $D, Q, \Delta, T, \epsilon = T_1, T_2, \dots, \epsilon$ .
1:  $\rho = \text{Lap}(\frac{\Delta}{2\epsilon})$ , count = 0
2: for each query  $q_i \in Q$  do
3:    $v_i = \text{Lap}(\frac{\Delta}{2\epsilon})$ 
4:   if  $q_i(D) + v_i \geq T_i + \rho$  then
5:     Output  $a_i = T$ 
6:     count = count + 1, Abort if count  $\geq \epsilon$ .
7:   else
8:     Output  $a_i = \pm 1$ 
9:   end if
10: end for
```

Algorithm 2 SVT in Dwork and Roth 2014 [8]

```
Input:  $D, Q, \Delta, T, \epsilon$ .
1:  $\rho = \text{Lap}(\frac{\Delta}{2\epsilon})$ , count = 0
2: for each query  $q_i \in Q$  do
3:    $v_i = \text{Lap}(\frac{\Delta}{2\epsilon})$ 
4:   if  $q_i(D) + v_i \geq T + \rho$  then
5:     Output  $a_i = T$ ,  $\rho = \text{Lap}(\frac{\Delta}{2\epsilon})$ 
6:     count = count + 1, Abort if count  $\geq \epsilon$ .
7:   else
8:     Output  $a_i = \pm 1$ 
9:   end if
10: end for
```

Algorithm 3 SVT in Roth's 2011 Lecture Notes [15]

```
Input:  $D, Q, \Delta, T, \epsilon$ .
1:  $\rho = \text{Lap}(\frac{\Delta}{2\epsilon})$ , count = 0
2: for each query  $q_i \in Q$  do
3:    $v_i = \text{Lap}(\frac{\Delta}{2\epsilon})$ 
4:   if  $q_i(D) + v_i \geq T_i + \rho$  then
5:     Output  $a_i = q_i(D) + v_i$ 
6:     count = count + 1, Abort if count  $\geq \epsilon$ .
7:   else
8:     Output  $a_i = \pm 1$ 
9:   end if
10: end for
```

Algorithm 5 SVT in Strohmann et al. 2014 [18]

```
Input:  $D, Q, \Delta, T$ .
1:  $\rho = \text{Lap}(\frac{\Delta}{2\epsilon})$ 
2: for each query  $q_i \in Q$  do
3:    $v_i = 0$ 
4:   if  $q_i(D) + v_i \geq T_i + \rho$  then
5:     Output  $a_i = T$ 
6:   else
7:     Output  $a_i = \pm 1$ 
8:   end if
9: end if
10: end for
```

Algorithm 4 SVT in Lee and Clifton 2014 [13]

```
Input:  $D, Q, \Delta, T, \epsilon$ .
1:  $\rho = \text{Lap}(\frac{\Delta}{2\epsilon})$ , count = 0
2: for each query  $q_i \in Q$  do
3:    $v_i = \text{Lap}(\frac{\Delta}{2\epsilon})$ 
4:   if  $q_i(D) + v_i \geq T + \rho$  then
5:     Output  $a_i = T$ 
6:     count = count + 1, Abort if count  $\geq \epsilon$ .
7:   else
8:     Output  $a_i = \pm 1$ 
9:   end if
10: end for
```

Algorithm 6 SVT in Chen et al. 2015 [11]

```
Input:  $D, Q, \Delta, T = T_1, T_2, \dots$ .
1:  $\rho = \text{Lap}(\frac{\Delta}{2\epsilon})$ 
2: for each query  $q_i \in Q$  do
3:    $v_i = \text{Lap}(\frac{\Delta}{2\epsilon})$ 
4:   if  $q_i(D) + v_i \geq T_i + \rho$  then
5:     Output  $a_i = T$ 
6:   else
7:     Output  $a_i = \pm 1$ 
8:   end if
9: end if
10: end for
```

	Alg. 1	Alg. 2	Alg. 3	Alg. 4	Alg. 5	Alg. 6
Scale of threshold noise ρ	$\frac{\Delta}{2\epsilon}$ or ϵ	$\frac{\Delta}{2\epsilon}$	$\frac{\Delta}{2\epsilon}$ or ϵ	$\frac{\Delta}{2\epsilon}$ or ϵ	$\frac{\Delta}{2\epsilon}$ or ϵ	$\frac{\Delta}{2\epsilon}$ or ϵ
Reset ρ after each output of T	No	Yes	No	No	No	No
Scale of query noise v_i	$\frac{\Delta}{2\epsilon}$ or ϵ	$\frac{\Delta}{2\epsilon}$	$\frac{\Delta}{2\epsilon}$ or ϵ	$\frac{\Delta}{2\epsilon}$ or ϵ	0	$\frac{\Delta}{2\epsilon}$ or ϵ
Outstripping $q_i + v_i$ instead of T	No	No	Yes	No	No	No
Stop after outputting $\epsilon + 1$	Yes	Yes	Yes	Yes	No	No
Privacy Property	ϵ -DP	ϵ -DP	∞ -DP ¹⁷	$(\frac{1}{2\epsilon}, \frac{1}{2\epsilon})$ -DP	∞ -DP	∞ -DP

Figure 2: Differences among Algorithms 1-6.

¹⁷ Algorithms 1 and 6 use a sequence of thresholds $T = T_1, T_2, \dots$, allowing different thresholds for different queries. The other algorithms use the same threshold T for all queries. We point out that this difference is mostly syntactical. In fact, having an SVT where the threshold always equals 0 suffices. Given a sequence of queries q_1, q_2, \dots , and a sequence of thresholds $T = T_1, T_2, \dots$, we can define a new sequence of queries $r_i = q_i - T_i$, and apply the SVT to r_i using 0 as the threshold to obtain the same result. In this paper, we decide to use thresholds to be consistent with the existing papers.

¹⁸ ∞ -DP means that an algorithm doesn't satisfy ϵ -DP for any finite privacy budget ϵ .

Probabilistic couplings

- ▶ Let μ_1 and μ_2 are distributions over A
- ▶ Let μ is a distribution over $A \times A$
- ▶ μ is a coupling of (μ_1, μ_2) if $\pi_1(\mu) = \mu_1$ and $\pi_2(\mu) = \mu_2$

Examples of couplings

- ▶ product distribution and optimal coupling
- ▶ $\mathcal{U}_{\{(1,1),(-1,-1)\}}$ is a coupling of $(\mathcal{U}_{\{1,-1\}}, \mathcal{U}_{\{1,-1\}})$
- ▶ $\mathcal{U}_{\{(1,-1),(-1,1)\}}$ is a coupling of $(\mathcal{U}_{\{1,-1\}}, \mathcal{U}_{\{1,-1\}})$

Benefits

- ▶ (Almost) no probabilistic reasoning
- ▶ Mechanizable
- ▶ Many proof techniques

Approximate couplings

- ▶ Let μ_1 and μ_2 be distributions over A
- ▶ Let μ_L and μ_R be distributions over $A \times A$
- ▶ (μ_L, μ_R) is an (ϵ, δ) -coupling of (μ_1, μ_2) if
 - ▶ $\pi_1(\mu_L) = \mu_1$ and $\pi_2(\mu_R) = \mu_2$
 - ▶ $\Delta_\epsilon(\mu_1, \mu_2) \leq \delta$

Benefits

- ▶ (Almost) no probabilistic reasoning
- ▶ Mechanizable
- ▶ Composition theorems extend
- ▶ New proof techniques
- ▶ Extend to f -divergences



Approximate probabilistic Relational Hoare Logic

- ▶ Judgment

$$\models \{P\} c_1 \sim_{\epsilon, \delta} c_2 \{Q\}$$

- ▶ Validity

$$\forall m_1, m_2. (m_1, m_2) \models P \implies (\llbracket c_1 \rrbracket m_1, \llbracket c_2 \rrbracket m_2) \models Q^{\#(\epsilon, \delta)}$$

- ▶ c is (ϵ, δ) -differentially private wrt Φ iff

$$\models \{\Phi\} c \sim_{\epsilon, \delta} c \{\equiv\}$$

Proof rules

$$\frac{\begin{array}{l} \models \{P\} c_1 \sim_{\epsilon, \delta} c_2 \{Q\} \\ \models \{Q\} c'_1 \sim_{\epsilon', \delta'} c'_2 \{R\} \end{array}}{\models \{P\} c_1; c'_1 \sim_{\epsilon+\epsilon', \delta+\delta'} c_2; c'_2 \{R\}}$$

$$\frac{\begin{array}{l} \models \{P \wedge e(1)\} c_1 \sim_{\epsilon, \delta} c \{Q\} \\ \models \{P \wedge \neg e(1)\} c_2 \sim_{\epsilon, \delta} c \{Q\} \end{array}}{\models \{P\} \text{ if } e \text{ then } c_1 \text{ else } c_2 \sim_{\epsilon, \delta} c \{Q\}}$$

$$\frac{\begin{array}{l} \models \{P \wedge e(1)\} c_1 \sim_{\epsilon, \delta} c'_1 \{Q\} \\ \models \{P \wedge \neg e(1)\} c_2 \sim_{\epsilon, \delta} c'_2 \{Q\} \\ P \rightarrow e(1) = e'(2) \end{array}}{\models \{P\} \text{ if } e \text{ then } c_1 \text{ else } c_2 \sim_{\epsilon, \delta} \text{ if } e' \text{ then } c'_1 \text{ else } c'_2 \{Q\}}$$

Proof principles for Laplace mechanism

Making different things look equal

$$\frac{\Phi \triangleq |e_1\langle 1 \rangle - e_2\langle 2 \rangle| \leq k'}{\vDash \{\Phi\} \quad y_1 \stackrel{\$}{\leftarrow} \text{Lap}_\epsilon(e_1) \sim_{k', \epsilon, 0} y_2 \stackrel{\$}{\leftarrow} \text{Lap}_\epsilon(e_2) \quad \{y_1\langle 1 \rangle = y_2\langle 2 \rangle\}}$$

Making equal things look different

$$\frac{\Phi \triangleq e_1\langle 1 \rangle = e_2\langle 2 \rangle}{\vDash \{\Phi\} \quad y_1 \stackrel{\$}{\leftarrow} \text{Lap}_\epsilon(e_1) \sim_{k, \epsilon, 0} y_2 \stackrel{\$}{\leftarrow} \text{Lap}_\epsilon(e_2) \quad \{y_1\langle 1 \rangle + k = y_2\langle 2 \rangle\}}$$

Keeping things the same, at no cost

$$\frac{y_1 \notin FV(e_1) \quad y_2 \notin FV(e_2) \quad \Psi \triangleq y_1\langle 1 \rangle - y_2\langle 2 \rangle = e_1\langle 1 \rangle - e_2\langle 2 \rangle}{\vDash \{\Psi\} \quad y_1 \stackrel{\$}{\leftarrow} \text{Lap}_\epsilon(e_1) \sim_{0, 0} y_2 \stackrel{\$}{\leftarrow} \text{Lap}_\epsilon(e_2) \quad \{\Psi\}}$$

Pointwise equality

$$\frac{\forall i. \vDash \{\Phi\} \quad c_1 \sim_{\epsilon, 0} c_2 \quad \{x\langle 1 \rangle = i \Rightarrow x\langle 2 \rangle = i\}}{\vDash \{\Phi\} \quad c_1 \sim_{\epsilon, 0} c_2 \quad \{x\langle 1 \rangle = x\langle 2 \rangle\}}$$

Coupling proof of sparse vector

Case $b = +\infty$ and $M = 1$

- ▶ (Cost ϵ): set $A\langle 1 \rangle + 1 = A\langle 2 \rangle$
- ▶ By pointwise equality, must prove for all k

$$I\langle 1 \rangle = k \Rightarrow I\langle 2 \rangle = k$$

- ▶ (Cost ϵ) critical iteration $i = k$: set $S\langle 1 \rangle + 1 = S\langle 2 \rangle$, and hence left test succeeds iff right test succeeds
- ▶ (Cost 0) iterations $i < k$: by sensitivity, $|S\langle 1 \rangle - S\langle 2 \rangle| \leq 1$, therefore right test succeeds implies left test succeeds
- ▶ (Cost 0) iterations $i > k$: similar

General case

- ▶ Use new optimal subset coupling for critical iterations
- ▶ Use accuracy to ensure that noisy intervals are non-empty

Accuracy via union bound logics

- ▶ Judgment $\models_{\beta} \{\Phi\} c \{\Psi\}$
- ▶ Validity: for every m , $m \models \Phi$ implies $\Pr_{[c](m)}[\neg\Psi] \leq \beta$

Selected rules

$$\frac{\models_{\beta} \{\Phi\} c \{\Phi'\} \quad \models_{\beta'} \{\Phi'\} c' \{\Phi''\}}{\models_{\beta+\beta'} \{\Phi\} c; c' \{\Phi''\}}$$

$$\frac{\models_{\beta} \{\Phi \wedge e\} c \{\Psi\} \quad \models_{\beta} \{\Phi \wedge \neg e\} c' \{\Psi\}}{\models_{\beta} \{\Phi\} \text{ if } e \text{ then } c \text{ else } c' \{\Psi\}}$$

$$\frac{\beta \in (0, 1) \quad \gamma = \frac{1}{\epsilon} \log \left(\frac{1}{\beta} \right)}{\models_{\beta} \{\top\} x \stackrel{s}{\leftarrow} \text{Lap}_{\epsilon}(e) \{|x - e| \leq \gamma\}}$$

Further work

- ▶ probabilistic Hoare logic:

$$\models \{\eta\} c \{\eta'\}$$

⇒ better accuracy and randomized algorithms

- ▶ probabilistic relational Hoare logic ($\epsilon = \delta = 0$):

$$\models \{\Psi\} c_1 \sim c_2 \{\Phi\}$$

⇒ cryptography and mechanism design

- ▶ proof-relevant probabilistic relational Hoare logic:

$$\models \{\Psi\} c_1 \sim c_2 \{\Phi\} \rightsquigarrow c$$

⇒ convergence of probabilistic processes

Formalization

```
+ while true (N - 1); last by auto => /#,
move>> z; wp; md pred; call adv_ll; auto; smt(lap_ll),
+ while true (N - 1); last by auto => /#,
move>> z; wp; md pred; call adv_ll; auto; smt(lap_ll),
+ by move>> /#,
move => [L GA] => /#,
case (size L = M),
conseq <- [big1 pred (fun x => int) =>
  if seen L (N - 1 - x) then 2hr * eps else @hr 0 N) & @hr ]>,
move>> s1 s2 /> _ lt LM; rewrite addrk -big_mkcond,
rewrite (partition_big (fun i => N - 1 - i) _ pred _ _ (undup L)) /#,
+ by rewrite undup_undup,
+ by move>> /> y; rewrite seen_undup,
apply(ler_trans (big pred (fun _ => 2hr * eps) (undup L)))>,
rewrite (Bigreal.sum_const mulrAC (mulrC (count _ )hr)),
rewrite ler_wgeu12 (mulr_ge0 // count_pred le_fromint,
+ by apply(ler_order. ler_trans _ _ (size_undup _)),
apply(Bigreal. ler_sum> 1 _ />); rewrite big_mkcond />,
case: (0 <= N - 1 - i < N) => h; last first,
rewrite -big_mkcond big_seq_cond big_pre0 (mulr_ge0 //,
+ by move>> ); rewrite seen_range /#,
rewrite (bigD1 _ (N - 1 - i)) (seen_range /range_undup //),
+ by rewrite big1 /!# /#; case: (L /> _ ) => /#; rewrite mulr_ge0,
while [ (fun x => if (seen L (N - 1 - x)) then 2hr * eps else @hr )
& (fun _ => @hr) ] N (N - 1 - 1) (
  adj d(1) d(2) /\ =i(1) /\ t@C(1) = t@C(2) + 1 /\
  (suffix (2) L => =(l, glob AJ) => //; first 4 by smt(ge0_eps),
+ by rewrite sumr_const intrmul,
+ move => v,
case (suffix (1(2)::(2)) L),
conseq <- [ (2hr * eps) & @hr ]>,
smt(suffix_cons),
seq 1 1: (adj d(1) d(2) /\
  =i(1, l, q, glob AJ) /\
  t@C(1) = t@C(2) + 1 /\ suffix (1(2)::(2)) L /\ v = N - i(1) - 1),
toequiv,
call ( _ : true); skip; smt(suffix1 suffix_cons),
seq 1 1: (adj d(1) d(2) /\
  =i(1, l, q, glob AJ) /\
  t@C(1) = t@C(2) + 1 /\ suffix (1(2)::(2)) L /\ v = N - i(1) - 1 /\
  suffix (1(2) :: (2) L /\ =i(1)-w(2)-1) <- (2hr * eps) & @hr ]>,
lap (-1 2 => />); smt(cons_sens),
toequiv auto,
smt,
toequiv,
smt(ge0_eps),
case (suffix (2) L),
seq 1 1: (adj d(1) d(2) /\
  =i(1, q) /\
  t@C(1) = t@C(2) + 1 /\
  suffix (2) L /\ =i(1, glob AJ) /\
  i(1) < N /\ v = N - i(1) - 1 /\
  | suffix (1(2) :: (2) L),
call ( _ : true); skip; smt(suffix1 suffix_cons),
wp,
exists* (eval0 q d(1)); elims => e1;
|- isv.ec 514 L330 Gitmaster (EasyCrypt script Scripting )
tool-bar goto
```

Current goal (remaining: 2)

Type variables: <none>

L: int List
GA: (glob AJ)

```
forall (k : int),
sequiv[[if seen L (N - 1 - k) then 2hr * eps else @hr & @hr] M(A).main.q <=
  A.adv(M(A).main.l(hr)); ...; M(A).main.i <-
  M(A).main.i(hr) = 1 - M(A).main.q <= A.adv(M(A).main.l(hr));
  ...; M(A).main.i <- M(A).main.i(hr) = 1 :
  (adj d(1) d(2) /\
  =i(1) /\ t@C(1) = t@C(2) + 1 /\ (suffix (2) L => =(l, glob AJ) /\
  i(1) < N /\ i(2) < N /\ k = N - i(1) - 1 /\ N - i(1) - 1 <= N =>
  (adj d(1) d(2) /\
  =i(1) /\ t@C(1) = t@C(2) + 1 /\ (suffix (2) L => =(l, glob AJ) /\
  i(1) < N - i(2) < N /\ N - i(1) - 1 < k]
```

U:\> -goals= All L1 (EasyCrypt goals)

+ >> Copyright (c) - 2012-2015 - IMDEA Software Institute and INRIA

>>> Distributed under the terms of the CeCILL-C license

>>> Standard Library (theories/na/ec):

>>> Distributed under the terms of the CeCILL-B license

>>>

>>> GIT hash: cb24e5da04ae7da56d1359d74b97ed6fe0669

U:\> -response= All L12 (EasyCrypt response)

EasyCrypt

- ▶ Interactive proof assistant for probabilistic programs
- ▶ Relational and non-relational program logics
- ▶ Back-end to SMT solvers and CAS
- ▶ libraries of common proof techniques (hybrid arguments, eager sampling, independent from adversary's view. . .)

Case studies

- ▶ encryption, signatures, hash designs, key exchange protocols, zero knowledge protocols, multi-party computation, verifiable computation. . .
- ▶ (computational) differential privacy
- ▶ mechanism design

Conclusion

- ▶ Fine-grained control key to “advanced” program verification
- ▶ Probabilistic RHL naturally models probabilistic couplings
- ▶ New perspectives on differential privacy and cryptography (beyond high assurance)
- ▶ Privacy and cryptography are great application domains (verified compilers, static analysis, synthesis. . .)