

# CS 591: Formal Methods in Security and Privacy

## Differential Privacy

Marco Gaboardi  
gaboardi@bu.edu

Alley Stoughton  
stough@bu.edu

# Feedback

Please fill out the (late) mid-semester evaluation.

# Recording

This is a reminder that we will record the class and we will post the link on Piazza.

This is also a reminder to myself to start recording!

From the previous classes

# A more realistic example

```
StreamCipher(m : private msg[n]) : public msg[n]  
  pkey := $ PRG(Uniform({0,1}k));  
  cipher := msg xor pkey;  
  return cipher
```

# How can we prove this secure?

```
OneTimePad(m : private msg[n])  
    : public msg[n]  
key := $ Uniform({0,1}n);  
cipher := msg xor key;  
return cipher
```

~

```
StreamCipher(m : private msg[n])  
    : public msg[n]  
pkey := $ PRG(Uniform({0,1}k));  
cipher := msg xor pkey;  
return cipher
```

# How can we prove this secure?

```
OneTimePad(m : private msg[n])  
    : public msg[n]  
key := $ Uniform({0,1}n);  
cipher := msg xor key;  
return cipher
```

m

~

```
StreamCipher(m : private msg[n])  
    : public msg[n]  
pkey := $ PRG(Uniform({0,1}k));  
cipher := msg xor pkey;  
return cipher
```

m

# How can we prove this secure?

```
OneTimePad(m : private msg[n])  
  : public msg[n]  
  key := $ Uniform({0,1}^n);  
  cipher := msg xor key;  
  return cipher
```

~

```
StreamCipher(m : private msg[n])  
  : public msg[n]  
  pkey := $ PRG(Uniform({0,1}^k));  
  cipher := msg xor pkey;  
  return cipher
```

m  
↓  
m ⊕ k

m



# How can we prove this secure?

```
OneTimePad(m : private msg[n])  
    : public msg[n]  
key := $ Uniform({0,1}^n);  
cipher := msg xor key;  
return cipher
```

~

```
StreamCipher(m : private msg[n])  
    : public msg[n]  
pkey := $ PRG(Uniform({0,1}^k));  
cipher := msg xor pkey;  
return cipher
```

m  
↓  
 $m \oplus k$

m  
↓  
 $m \oplus pk$

# How can we prove this secure?

```
OneTimePad(m : private msg[n])  
  : public msg[n]  
  key := $ Uniform({0,1}^n);  
  cipher := msg xor key;  
  return cipher
```

~

```
StreamCipher(m : private msg[n])  
  : public msg[n]  
  pkey := $ PRG(Uniform({0,1}^k));  
  cipher := msg xor pkey;  
  return cipher
```

$$\begin{array}{ccc} m & & m \\ \downarrow & & \downarrow \\ \Delta(m \oplus k) & , & m \oplus p k \leq \delta \end{array}$$

# Approximate Probabilistic Relational Hoare Logic

Indistinguishability  
parameter

Precondition  
(a logical formula)

$$\vdash_{\delta} c_1 \sim c_2 : P \Rightarrow Q$$

Probabilistic  
Program

Probabilistic  
Program

Postcondition  
(a logical formula)

# Validity of approximate Probabilistic Hoare judgments

We say that the quadruple  $\vdash_{\delta} c_1 \sim c_2 : P \Rightarrow Q$  is **valid** if and only if for every pair of memories  $m_1, m_2$  such that  $P(m_1, m_2)$  we have:

$\{c_1\}_{m_1} = \mu_1$  and  $\{c_2\}_{m_2} = \mu_2$  implies

$Q_{\delta^*}(\mu_1, \mu_2)$ .

# R- $\delta$ -Coupling

Given two distributions  $\mu_1 \in D(A)$ , and  $\mu_2 \in D(B)$ , we have an R- $\delta$ -coupling between them, for  $R \subseteq A \times B$  and  $0 \leq \delta \leq 1$ , if there are two joint distributions  $\mu_L, \mu_R \in D(A \times B)$  such that:

- 1)  $\pi_1(\mu_L) = \mu_1$  and  $\pi_2(\mu_R) = \mu_2$ ,
- 2) the support of  $\mu_L$  and  $\mu_R$  is contained in  $R$ .  
That is, if  $\mu_L(a, b) > 0$ , then  $(a, b) \in R$ ,  
and if  $\mu_R(a, b) > 0$ , then  $(a, b) \in R$ .
- 3)  $\Delta(\mu_L, \mu_R) \leq \delta$

# Probabilistic Relational Hoare Logic

## Skip

---

$$\vdash_0 \text{skip} \sim \text{skip} : P \Rightarrow P$$

# Probabilistic Relational Hoare Logic

## Composition

$$\vdash \delta_1 C_1 \sim C_2 : P \Rightarrow R \quad \vdash \delta_2 C_1' \sim C_2' : R \Rightarrow S$$

---

$$\vdash \delta_1 + \delta_2 C_1 ; C_1' \sim C_2 ; C_2' : P \Rightarrow S$$

# Probabilistic Relational Hoare Logic

## A specific rule for PRG

---

$$\vdash_{2^{-n}} \begin{array}{l} x_1 := \$ \text{Uniform}(\{0, 1\}^n) \sim \\ x_2 := \$ \text{PRG}(\text{Uniform}(\{0, 1\}^k)) \\ : \text{True} \Rightarrow x_1\langle 1 \rangle = x_2\langle 2 \rangle \end{array}$$



# How can we prove this secure?

```
OneTimePad(m : private msg[n])  
  : public msg[n]  
  key := $ Uniform({0,1}^n);  
  cipher := msg xor key;  
  return cipher
```

~

```
StreamCipher(m : private msg[n])  
  : public msg[n]  
  pkey := $ PRG(Uniform({0,1}^k));  
  cipher := msg xor pkey;  
  return cipher
```

We can apply the PRG rule, the composition rule, and the assignment rule and prove:

```
⊢2-n OneTimePad ~ StreamCipher  
: m<1> = m<2> ⇒ c<1> = c<2>
```

# Differential Privacy

# Data



**Aol.**



Google

# Releasing the mean of Some Data

```
Mean (d : private data) : public real
  i:=0;
  s:=0;
  while (i<size(d))
    s:=s + d[i]
    i:=i+1;
  return (s/i)
```


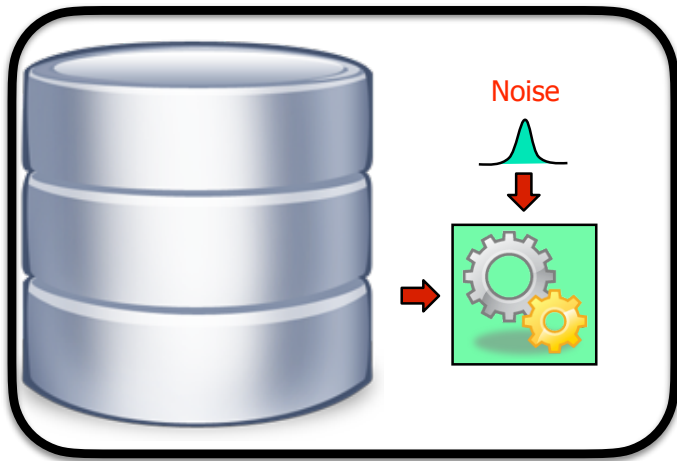
# Privacy-preserving data analysis?

We want to release some information to a data analyst and protect the privacy of the individuals contributing their data.





# Privacy-preserving data analysis?

We want to release some information to a data analyst and protect the privacy of the individuals contributing their data.



# Data analyst




# Fundamental Law of Information Reconstruction

The release of **too many** overly **accurate** statistics permits reconstruction attacks.







# Reconstruction attack





Attacker





# Reconstruction attack





# Reconstruction attack



# Reconstruction attack



# Reconstruction attack



# Reconstruction attack



We say that the attacker **wins** if

$$d(\text{D}, \text{D}') \sim 0$$

# Reconstruction attack




We say that the attacker **wins** if

$$d(\text{D}, \text{D}') \sim 0$$

In this class case we can use Hamming distance

# Privacy vs Utility





# Quantitative notions of Privacy

- The impossibility results discussed above suggest a quantitative notion of privacy,
- a notion where the privacy loss depends on the number of queries that are allowed,
- and on the accuracy with which we answer them.

Differential privacy:  
understanding the mathematical and  
computational meaning of this trade-  
off.

[Dwork, McSherry, Nissim, **Smith**, TCC06]

# Privacy-preserving data analysis?

- The analyst knows no more about me after the analysis than what she knew before the analysis.


# Privacy-preserving data analysis?

- The analyst knows no more about me after the analysis than what she knew before the analysis.



# Privacy-preserving data analysis?

- The analyst knows no more about me after the analysis than what she knew before the analysis.



# Privacy-preserving data analysis?

Prior Knowledge

~

Posterior Knowledge


Privacy-preserving data analysis?

# Privacy-preserving data analysis?

**Question:** What is the problem with this requirement?



# Privacy-preserving data analysis?



If nothing can be learned about an individual, then nothing at all can be learned at all!

[DworkNaor10]

# Privacy-preserving data analysis?

- The analyst learn **almost the same** about me after the analysis as what she would have learnt if I **didn't contribute my data**.


# Privacy-preserving data analysis?

- The analyst learn **almost the same** about me after the analysis as what she would have learnt if I **didn't contribute my data**.




# Privacy-preserving data analysis?

- The analyst learn **almost the same** about me after the analysis as what she would have learnt if I **didn't contribute my data**.




# Privacy-preserving data analysis?

- The analyst learn **almost the same** about me after the analysis as what she would have learnt if I **didn't contribute my data**.



# Privacy-preserving data analysis?

- The analyst learn **almost the same** about me after the analysis as what she would have learnt if I **didn't contribute my data**.



# Adjacent databases

- We can formalize the concept of contributing my data or not in terms of a notion of distance between datasets.
- Given two datasets  $D, D' \in DB$ , their distance is defined as:

$$D \Delta D' = |\{k \leq n \mid D(k) \neq D'(k)\}|$$

- We will call two datasets adjacent when  $D \Delta D' = 1$  and we will write  $D \sim D'$ .

# Privacy Loss

In general we can think about the following quantity as the **privacy loss** incurred by observing  $r$  on the databases  $b$  and  $b'$ .

$$L_{b,b'}(r) = \log \frac{\Pr[Q(b)=r]}{\Pr[Q(b')=r]}$$





# $(\epsilon, \delta)$ -Differential Privacy

## Definition


Given  $\epsilon, \delta \geq 0$ , a probabilistic query  $Q: X^n \rightarrow R$  is  $(\epsilon, \delta)$ -differentially private iff  
for all adjacent database  $b_1, b_2$  and for every  $S \subseteq R$ :

$$\Pr[Q(b_1) \in S] \leq \exp(\epsilon) \Pr[Q(b_2) \in S] + \delta$$


# Differential Privacy

$Q : \mathcal{D} \Rightarrow \mathcal{R}$  probabilistic

$Q(\mathcal{D}_x)$




$Q(\mathcal{D}_y)$



# Differential Privacy

$d(Q(\text{bu}\{x\}), Q(\text{bu}\{y\})) \leq \epsilon$  with probability  $1-\delta$



# $(\epsilon, \delta)$ -Differential Privacy

