

COS433/Math 473: Cryptography

Mark Zhandry

Princeton University

Spring 2020

Previously on COS 433...

Perfect Security for Multiple Messages

Definition: A stateless scheme **(Enc, Dec)** has **perfect secrecy for n messages** if, for any two sequences of messages $(m_0^{(i)})_{i \in [d]}$, $(m_1^{(i)})_{i \in [d]} \in M^d$

$$\left(\text{Enc}(K, m_0^{(i)}) \right)_{i \in [d]} \stackrel{d}{=} \left(\text{Enc}(K, m_1^{(i)}) \right)_{i \in [d]}$$

Notation: $(f(i))_{i \in [d]} = (f(1), f(2), \dots, f(n))$

Randomized Encryption

Syntax:

- Key space \mathbf{K} (usually $\{0,1\}^\lambda$)
- Message space \mathbf{M} (usually $\{0,1\}^n$)
- Ciphertext space \mathbf{C} (usually $\{0,1\}^m$)
- **Enc:** $\mathbf{K} \times \mathbf{M} \rightarrow \mathbf{C}$ (potentially probabilistic)
- **Dec:** $\mathbf{K} \times \mathbf{C} \rightarrow \mathbf{M}$ (usually deterministic)

Correctness:

- For all $k \in \mathbf{K}$, $m \in \mathbf{M}$,
$$\Pr[\text{Dec}(k, \text{Enc}(k,m)) = m] = 1$$

Theorem: No stateless *randomized* encryption scheme can have perfect security for multiple messages

Today: Relaxing
Perfect Secrecy

What do we do now?

Tolerate tiny probability of distinguishing

- If $\Pr[\mathbf{c}^{(0)} = \mathbf{c}^{(1)}] = 2^{-128}$, in reality never going to happen

How Small Is Ok?

Practice:

- Something unlikely to happen in lifetime of data/person/civilization/universe
- Typically something like **2^{-80}** , **2^{-128}** , or maybe **2^{-258}**
 - Being struck by lightning twice: **2^{-23}**
 - Winning the lottery: **2^{-26}**
 - World-ending asteroid while on this slide: **2^{-46}**

How Small Is Ok?

Theory:

- Maybe things will change as technology improves
- Want a more conceptual answer
- Absolute constants unsatisfactory
- Instead, use ``negligible'' functions

Negligible functions

Def: A function \mathbf{f} is **polynomial** if $\mathbf{f}(n) = O(n^c)$ for some constant \mathbf{c}

Def: A function \mathbf{g} is **super-polynomial** if, for every polynomial \mathbf{f} , $\mathbf{f}(n) = O(\mathbf{g}(n))$

Def: A function \mathbf{p} is **inverse polynomial** if $\mathbf{1/p}(n)$ is polynomial

Def: A function $\boldsymbol{\varepsilon}$ is **negligible** if, for every inverse polynomial \mathbf{p} , $\boldsymbol{\varepsilon}(n) = O(\mathbf{p}(n))$

(equivalently, $\mathbf{1/\varepsilon}$ is super-polynomial)

Examples

2^n super-polynomial

$n^{-n/7}$ negligible

$3^{-5 \log n}$ inverse polynomial

$1.5^{-\sqrt[3]{n}}$ negligible

$8^{\log^3 n}$ super-polynomial

$(\log n)/n$ inverse polynomial

Security Parameter λ

Additional input to system, dictates “security level”

Key, message, ciphertext size all **polynomial** in λ

Probability of adversary success is **negligible** in λ

Defining Encryption Again

Syntax:

- Key space \mathbf{K}_λ
- Message space \mathbf{M}_λ
- Ciphertext space \mathbf{C}_λ
- **Enc:** $\mathbf{K}_\lambda \times \mathbf{M}_\lambda \rightarrow \mathbf{C}_\lambda$ (potentially randomized)
- **Dec:** $\mathbf{K}_\lambda \times \mathbf{C}_\lambda \rightarrow \mathbf{M}_\lambda$

Correctness:

- $\log|\mathbf{K}_\lambda|, \log|\mathbf{M}_\lambda|, \log|\mathbf{C}_\lambda|$ polynomial in λ
- For all $\lambda, k \in \mathbf{K}_\lambda, m \in \mathbf{M}_\lambda$,
$$\Pr[\Pr[\text{Dec}(k, \text{Enc}(k,m)) = m] = 1] = 1$$

Statistical Distance

Given two distributions $\mathbf{D}_1, \mathbf{D}_2$ over a set \mathbf{X} , define

$$\Delta(\mathbf{D}_1, \mathbf{D}_2) = \frac{1}{2} \sum_{\mathbf{x}} | \Pr[\mathbf{D}_1 = \mathbf{x}] - \Pr[\mathbf{D}_2 = \mathbf{x}] |$$

Observations:

$$0 \leq \Delta(\mathbf{D}_1, \mathbf{D}_2) \leq 1$$

$$\Delta(\mathbf{D}_1, \mathbf{D}_2) = 0 \iff \mathbf{D}_1 \stackrel{d}{=} \mathbf{D}_2$$

$$\Delta(\mathbf{D}_1, \mathbf{D}_2) \leq \Delta(\mathbf{D}_1, \mathbf{D}_3) + \Delta(\mathbf{D}_3, \mathbf{D}_2)$$

(Δ is a metric)

Another View of Statistical Distance

Theorem: $\Delta(\mathcal{D}_1, \mathcal{D}_2) \geq \varepsilon$ iff \exists (potentially randomized) \mathbf{A} s.t.

$$\left| \Pr[\mathbf{A}(\mathcal{D}_1) = 1] - \Pr[\mathbf{A}(\mathcal{D}_2) = 1] \right| \geq \varepsilon$$

Terminology: for any \mathbf{A} ,
 $\left| \Pr[\mathbf{A}(\mathcal{D}_1) = 1] - \Pr[\mathbf{A}(\mathcal{D}_2) = 1] \right|$
is called the “advantage” of \mathbf{A} in
distinguishing \mathcal{D}_1 and \mathcal{D}_2

Another View of Statistical Distance

Theorem: $\Delta(\mathcal{D}_1, \mathcal{D}_2) \geq \varepsilon$ iff \exists (potentially randomized) \mathbf{A} s.t.

$$\left| \Pr[\mathbf{A}(\mathcal{D}_1) = 1] - \Pr[\mathbf{A}(\mathcal{D}_2) = 1] \right| \geq \varepsilon$$

To lower bound Δ , just need to show adversary \mathbf{A} with that advantage

Examples

D_1 = Uniform distribution over $\{0,1\}^n$

- $\Pr[D_1=x] = 2^{-n}$

D_2 = Uniform subject to even parity

- $\Pr[D_2=x] = 2^{-(n-1)}$ if x has even parity, 0 otherwise

$$\begin{aligned}\Delta(D_1, D_2) &= \frac{1}{2} \sum_{\text{even } x} |2^{-n} - 2^{-(n-1)}| \\ &\quad + \frac{1}{2} \sum_{\text{odd } x} |2^{-n} - 0| \\ &= \frac{1}{2} \sum_{\text{even } x} 2^{-n} + \frac{1}{2} \sum_{\text{odd } x} 2^{-n} \\ &= \frac{1}{2}\end{aligned}$$

Examples

$\mathcal{D}_1 =$ Uniform over $\{1, \dots, n\}$

$\mathcal{D}_2 =$ Uniform over $\{1, \dots, n+1\}$

$$\begin{aligned}\Delta(\mathcal{D}_1, \mathcal{D}_2) &= \frac{1}{2} \sum_{x=1}^n \left| \frac{1}{n} - \frac{1}{n+1} \right| \\ &\quad + \frac{1}{2} \left| 0 - \frac{1}{n+1} \right| \\ &= \frac{1}{2} \sum_{x=1}^n \frac{1}{n(n+1)} + \frac{1}{2} \frac{1}{n+1} \\ &= \frac{1}{2} \frac{1}{n+1} + \frac{1}{2} \frac{1}{n+1} = \frac{1}{n+1}\end{aligned}$$

Statistical Security (Concrete)

Definition: A scheme **(Enc, Dec)** has **ϵ -statistical secrecy for d messages** if \forall two sequences of messages $(m_0^{(i)})_{i \in [d]}$, $(m_1^{(i)})_{i \in [d]} \in M^d$

$$\Delta \left[\left(\text{Enc}(K, m_0^{(i)}) \right)_{i \in [d]}, \left(\text{Enc}(K, m_1^{(i)}) \right)_{i \in [d]} \right] < \epsilon$$

We will call such a scheme **(d, ϵ) statistically secure**

Statistical Security (Asymptotic)

Definition: A scheme **(Enc, Dec)** has **statistical secrecy for d messages** if \exists negligible ε such that \forall two sequences $(m_0^{(i)})_{i \in [d]}$, $(m_1^{(i)})_{i \in [d]} \in \mathcal{M}_\lambda^d$,

$$\Delta \left[\left(\text{Enc}(K_\lambda, m_0^{(i)}) \right)_{i \in [d]}, \left(\text{Enc}(K_\lambda, m_1^{(i)}) \right)_{i \in [d]} \right] < \varepsilon(\lambda)$$

We will call such a scheme **d -time statistically secure**

Stateless Encryption with Multiple Messages

Ex:

$$\mathbf{M = C = \mathbb{Z}_p \text{ (} p \text{ a prime of size } 2^\lambda, \lambda=128)}$$

$$\mathbf{K = \mathbb{Z}_p^* \times \mathbb{Z}_p}$$

$$\mathbf{Enc((a,b), m) = (am + b) \bmod p}$$

$$\mathbf{Dec((a,b), c) = (c-b)/a \bmod p}$$

Q: Is this statistically secure for two messages?

Example

Ex:

$$\mathbf{M} = \mathbb{Z}_p \text{ (} p \text{ a prime of size } 2^\lambda, \lambda=128 \text{)}$$

$$\mathbf{C} = \mathbb{Z}_p^2$$

$$\mathbf{K} = \mathbb{Z}_p^2$$

$$\text{Enc}((a,b), m) = (r, (ar+b) + m)$$

$$\text{Dec}((a,b), (r,c)) = c - (ar+b)$$

Random in \mathbb{Z}_p



Q: Is this statistically secure for two messages?

Proof of Example

Let \mathbf{D}_b be distribution of $(\mathbf{Enc}(k, m_b^{(i)}))_{i \in \{1,2\}}$

Let \mathbf{D}_b' be \mathbf{D}_b , but conditioned on $r_0 \neq r_1$

Fix $r_0 \neq r_1, m_0, m_1, c_0, c_1$

$$\Pr_{(a,b)}[ar_0 + b + m_0 = c_0, ar_1 + b + m_1 = c_1] = 1/p^2$$

$$\text{So } \mathbf{D}_0' \stackrel{d}{=} \mathbf{D}_1' \quad (\Delta(\mathbf{D}_0', \mathbf{D}_1') = 0)$$

Proof of Example

Lemma: $\Delta(\mathcal{D}_1, \mathcal{D}_2) \leq \frac{1}{2}\Pr[\text{bad}|\mathcal{D}_1] + \frac{1}{2}\Pr[\text{bad}|\mathcal{D}_2] + \Delta(\mathcal{D}_1', \mathcal{D}_2')$

Where:

- “**bad**” is some event
- $\Pr[\text{bad}|\mathcal{D}_b]$ is probability “**bad**” when sampling from \mathcal{D}_b
- \mathcal{D}_b' is \mathcal{D}_b , but conditioned on **not** “**bad**”

Proof of Lemma

$$\begin{aligned}\Delta(D_1, D_2) &= \frac{1}{2} \sum_x | \Pr[D_1=x] - \Pr[D_2=x] | \\ &= \frac{1}{2} \sum_{x:\text{bad}} | \Pr[D_1=x] - \Pr[D_2=x] | \\ &\quad + \frac{1}{2} \sum_{x:\text{good}} | \Pr[D_1=x] - \Pr[D_2=x] | \\ &\leq \frac{1}{2} \sum_{x:\text{bad}} | \Pr[D_1=x] | + \frac{1}{2} \sum_{x:\text{bad}} | \Pr[D_2=x] | \\ &\quad + \frac{1}{2} \sum_{x:\text{good}} | \Pr[D_1=x] - \Pr[D_2=x] | \\ &\leq \frac{1}{2} \Pr[\text{bad}|D_1] + \frac{1}{2} \Pr[\text{bad}|D_2] + \Delta(D_1', D_2')\end{aligned}$$

Proof of Example

Let \mathbf{D}_b be distribution of $(\mathbf{Enc}(k, m_b^{(i)}))_{i \in \{1,2\}}$

Let **bad** be when $r_0 = r_1$

Let \mathbf{D}_b' be \mathbf{D}_b , but conditioned on **not bad**

$$\Pr[\mathbf{bad} | \mathbf{D}_b] = 1/p$$

$$\Delta(\mathbf{D}_0', \mathbf{D}_1') = 0$$

Therefore, $\Delta(\mathbf{D}_0, \mathbf{D}_1) \leq 1/p \approx 2^{-\lambda}$

Summary so Far

Stateless encryption for multiple messages



But, key length grows with number of messages



And, key length grows with length of message



Limits of Statistical Security

Theorem: Suppose **(Enc,Dec)** has plaintext space $\mathbf{M} = \{0,1\}^n$ and key space $\mathbf{K} = \{0,1\}^t$. Moreover, assume it is **(d, 0.4999)**-secure. Then:

$$t \geq d n$$

In other words, the key must be at least as long as the total length of all messages encrypted

Proof Idea: Compression

Use an encryption protocol to build a compression protocol



m



$$m' \leftarrow \text{Comp}(m)$$


$$m \leftarrow \text{Decomp}(m')$$

$$\text{Goal: } |m'| < |m|$$

For Now: Easier Goal



m



$s \leftarrow \text{Setup}()$



$m \leftarrow \text{Decomp}(s, m')$

$m' \leftarrow \text{Comp}(s, m)$

Goal: $|m'| < |m|$

The Protocol

Let m_0 be some message in M

Setup():

- Choose random $k_0 \leftarrow K$
- Let $c_1 \leftarrow \text{Enc}(k_0, m_0), \dots, c_d \leftarrow \text{Enc}(k_0, m_0)$
- Output (c_1, \dots, c_d)


 In M^d

Comp($(c_1, \dots, c_d), (m_1, \dots, m_d)$):

- Find k, r_1, \dots, r_d such that $c_i = \text{Enc}(k, m_i; r_i) \quad \forall i$
- If no such values exist, abort
- Output k

The Protocol

Let \mathbf{m}_0 be some message in \mathbf{M}

Comp($(\mathbf{c}_1, \dots, \mathbf{c}_d)$, $(\mathbf{m}_1, \dots, \mathbf{m}_d)$):  In \mathbf{M}^d

- Find $\mathbf{k}, \mathbf{r}_1, \dots, \mathbf{r}_d$ such that $\mathbf{c}_i = \text{Enc}(\mathbf{k}, \mathbf{m}_i; \mathbf{r}_i) \quad \forall i$
- If no such values exist, abort
- Output \mathbf{k}

Decomp($(\mathbf{c}_1, \dots, \mathbf{c}_d)$, \mathbf{k}):

- Compute $\mathbf{m}_i = \text{Dec}(\mathbf{k}, \mathbf{c}_i)$
- Output $(\mathbf{m}_1, \dots, \mathbf{m}_d)$

Analysis of Protocol

If **Comp** succeeds, **Decomp** must succeed by correctness

- Since $c_i = \text{Enc}(k, m_i; r_i)$, $\text{Dec}(k, c_i)$ must give m_i

Therefore, must figure out when **Comp** succeeds

Claim: For any sequence of messages m_1, \dots, m_d , **Comp** succeeds with probability at least $1 - \epsilon$

(Probability over the randomness used by **Setup()**)

Claim: For any sequence of messages $\mathbf{m}_1, \dots, \mathbf{m}_d$, **Comp** succeeds with probability at least $1-\epsilon$

Proof:

- Suppose **Comp** succeeds with probability $1-p$ for messages $\mathbf{m}_1, \dots, \mathbf{m}_d$
- Let $\mathbf{A}(\mathbf{c}_1, \dots, \mathbf{c}_d)$ be the algorithm that runs $\mathbf{Comp}((\mathbf{c}_1, \dots, \mathbf{c}_d), (\mathbf{m}_1, \dots, \mathbf{m}_d))$ and outputs **1** if **Comp** succeeds
- If $\mathbf{c}_i = \mathbf{Enc}(\mathbf{k}_0, \mathbf{m}_i)$, then $\Pr[\mathbf{A}(\mathbf{c}_1, \dots, \mathbf{c}_d)=1] = 1$
- If $\mathbf{c}_i = \mathbf{Enc}(\mathbf{k}_0, \mathbf{m}_0)$, then $\Pr[\mathbf{A}(\mathbf{c}_1, \dots, \mathbf{c}_d)=1] = 1-p$
- By (d, ϵ) -statistical security of **Enc**, p must be $\leq \epsilon$

Claim: For any sequence of messages $\mathbf{m}_1, \dots, \mathbf{m}_d$,
Comp succeeds with probability at least $1-\epsilon$

Claim: For **a random** sequence of messages
 $\mathbf{m}_1, \dots, \mathbf{m}_d$, **Comp** succeeds with prob at least $1-\epsilon$

(Probability over the randomness used by **Setup()**
and the random choices of $\mathbf{m}_1, \dots, \mathbf{m}_d$)

Next step: Removing Setup

We know:

$$\Pr[\text{Comp succeeds: } \underset{m_i \leftarrow M}{(c_1, \dots, c_d) \leftarrow \text{Setup()}}] \geq 1 - \varepsilon$$


Therefore, there must exist *some* (c_1^*, \dots, c_d^*) such that

$$\Pr[\text{Comp succeeds: } m_i \leftarrow M] \geq 1 - \varepsilon$$

Define: $M' = \{(m_1, \dots, m_d): \text{Comp succeeds}\}$

- Note that $|M'| \geq (1 - \varepsilon) |M|^d$

The Protocol



Find k, r_1, \dots, r_d such that
 $c_i^* = \text{Enc}(k, m_i; r_i) \quad \forall i$

For each i ,
Let $m_i \leftarrow \text{Dec}(k, c_i^*)$
Output (m_1, \dots, m_d)

By previous analysis,

- Alice always successfully compresses
- Bob always successfully decompresses

Final Touches

Can compress messages in \mathbf{M}' into keys in \mathbf{K}


Therefore, it must be that $|\mathbf{M}'| \leq |\mathbf{K}|$

Meaning $t = \log |\mathbf{K}|$
 $\geq \log |\mathbf{M}'|$
 $\geq \log [(1-\varepsilon) |\mathbf{M}|^d]$
 $= d \log |\mathbf{M}| + \log [1-\varepsilon]$
 $= dn + \log [1-\varepsilon]$
 $\geq dn$ (as long as $\varepsilon < 1/2$)

Takeaway

If you don't want to physically exchange keys frequently, you cannot obtain statistical security

So, now what?



Computational Security

We are ok if adversary takes a really long time

Only considered attack for adversaries that don't take too long

How Long Is Ok?

Practice:

- Lifetime of data/person/civilization/universe
- Typically something like 2^{80} , 2^{128} , or maybe 2^{258}
 - Lifetime of universe in nanoseconds: 2^{58}
 - Number of atoms in known universe: 2^{265}

How Long Is Ok?

Theory:

- Maybe things will change as technology improves
- Want a more conceptual answer
- Absolute constants unsatisfactory
- Instead, consider an attack if time bounded by polynomial function


Brute Force Attacks

Simply try every key until find right one

If keys have length λ , 2^λ is upper bound on attack


Not always applicable. When?

Holiwudd Criptoe!



[TRANSLTR]'s three million processors would all work in parallel ... trying every new permutation as they went

Holiwudd Criptoe!



“What’s the longest you’ve ever seen TRANSLTR take to break a code?”

“About an hour, but it had a ridiculously long key—ten thousand bits”

Reminders

HW1 Due Feb 20th

Project 1 to be released hopefully this afternoon