CS 4770: Cryptography

CS 6750: Cryptography and Communication Security

Alina Oprea
Associate Professor, CCIS
Northeastern University

January 11 2018

CS 4770, CS 6750: Syllabus

- Symmetric-key primitives
  - Block ciphers, symmetric-key encryption
  - Pseudorandom functions and pseudorandom generators
  - MACs and authenticated encryption
- Hash functions
  - Integrity schemes
- Public-key cryptography
  - Public-key encryption and signatures
  - Key exchange
- Applications
  - Secure network communication, secure computation, crypto currencies

Textbook: Introduction to Modern Cryptography.
J. Katz and Y. Lindell

Policies

- Instructors
  - Alina Oprea
  - TA: Sourabh Marathe
- Schedule
  - Mon, Thu 11:45am – 1:25pm, Robinson 107
  - Office hours:
    - Alina: Thu 4:00 – 6:00 pm (ISEC 625)
    - Sourabh: Tue 2-3pm (ISEC 532)
- Your responsibilities
  - Please be on time and attend classes
  - Participate in interactive discussion
  - Submit assignments/ programming projects on time
- Late days for assignments
  - 5 total late days, after that loose 20% for every late day
  - Assignments are due at 11:59pm on the specified date
- Respect university code of conduct
  - No collaboration on homework / programming projects
  - http://www.northeastern.edu/osccr/academic-integrity-policy/

Grading

- Written problem assignments – 25%
  - 3-4 theoretical problem assignments based on studied material in class
- Programming projects – 20%
  - 3 programming projects
  - Language of your choice (Java, C/C++, Python)
  - In-person grading with instructor/TA
- Exams – 50%
  - Midterm – 25%
  - Final exam – 25%
- Class participation – 5%
  - Participate in class discussion and on Piazza

Review

- Historically cryptography used by military
  - All historical ciphers (shift, substitution, Vigenere) have been broken
  - If key space is small (shift cipher), can mount brute-force attack
  - Large key space doesn’t mean cipher is secure!
- Modern cryptography
  - Rooted in formal definitions and rigorous proofs based on computational assumptions
  - Enables a number of emerging applications

Outline

- Probability review
  - Events, union bound
  - Conditional probability, Bayes theorem
- Defining security for encryption
  - Several wrong approaches
- Perfect secrecy
  - Rigorous definition of security for encryption (Shannon 1949)
- One-time pad
  - Construction, proof and limitations

Probability review

Probability space and events

- Probability space:
  - Universe $\mathcal{U}$
  - Probability function: for all $\mathrm{u} \in \mathcal{U}$, assign $0 \le \Pr[u] \le 1$ such that $\sum_{u \in \mathcal{U}} \Pr[u] = 1$.
- Event is a set $A \subseteq \mathcal{U}$: $\Pr[A] = \sum_{x \in A} \Pr(x) \quad \in [0,1]$

note: $\Pr[\mathcal{U}]=1$

Example

- $\mathcal{U} = \{0,1\}^8$
- $A = \{$ all x in $\mathcal{U}$ such that $\mathrm{lsb}_2(x)=11$ $\} \subseteq \mathcal{U}$

for the uniform distribution on $\{0,1\}^8$:

$\Pr[A] = 1/4$

The union bound

- For events A_1 and A_2

$$\Pr[A_1 \cup A_2] \leq \Pr[A_1] + \Pr[A_2]$$

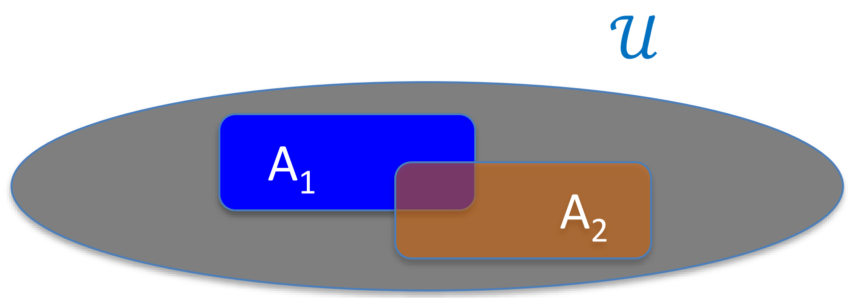

If $A_1 \cap A_2 = \Phi$, then $\Pr[A_1 \cup A_2] = \Pr[A_1] + \Pr[A_2]$

In general $\Pr[A_1 \cup A_2] = \Pr[A_1] + \Pr[A_2] - \Pr[A_1 \cap A_2]$

Example:

$A_1 = \{$ all x in $\{0,1\}^n$ s.t $lsb_2(x)=11$ $\}$; $A_2 = \{$ all x in $\{0,1\}^n$ s.t. $msb_2(x)=11$ $\}$

$\Pr[\, lsb_2(x)=11$ or $msb_2(x)=11\,] = \Pr[A_1 \cup A_2] \leq \frac{1}{4}+\frac{1}{4} = \frac{1}{2}$

Random Variables

Def: a random variable X is a function $X: U \longrightarrow V$

Example: $X: \{0,1\}^n \longrightarrow \{0,1\}$; $X(y) = \text{lsb}(y) \in \{0,1\}$

For the uniform distribution on U:

$\Pr[X=0] = 1/2$, $\Pr[X=1] = 1/2$

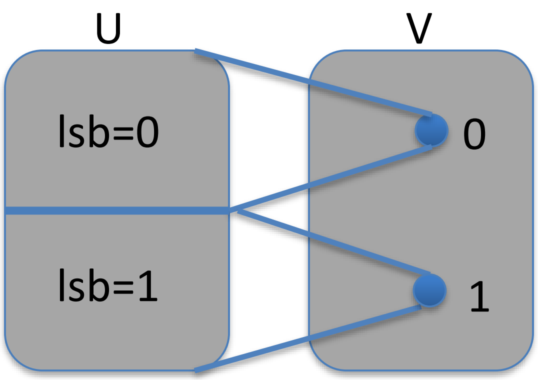

More generally:

Rand. var. X takes values in V and induces a distribution on V

The uniform random variable

Let U be some set, e.g. $U = \{0,1\}^n$

We write $r \xleftarrow{R} U$ to denote a **uniform random variable** over U

for all $u \in U$: $\Pr[\, r = u \,] = 1/|U|$

Randomized algorithms

- Deterministic algorithm: $y \longleftarrow A(m)$

- Randomized algorithm

  $y \longleftarrow A(\,m\,;\,r\,)$ where $r \xleftarrow{R} \{0,1\}^n$

  output is a random variable

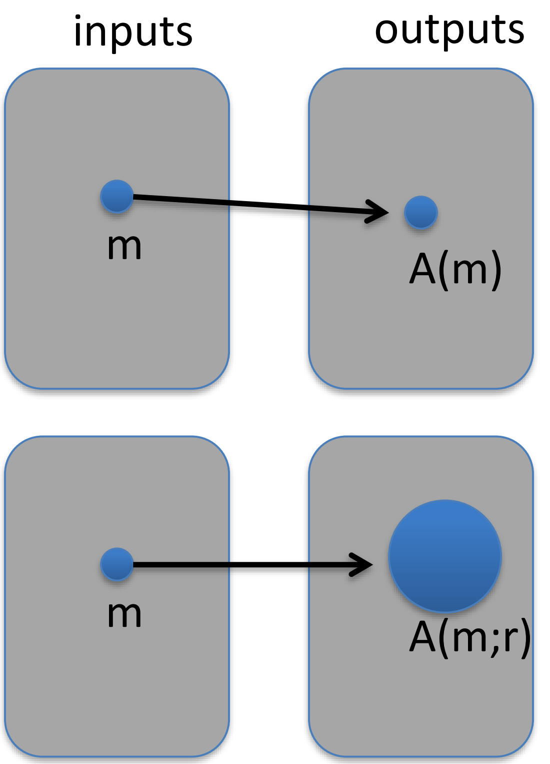

Example: $A(m\,;\,r) = m+r$

Independence

Def: Events A and B are **independent** if and only if

$$\Pr[\ A \text{ and } B\ ] = \Pr[A] \cdot \Pr[B]$$

Random variables X,Y taking values in V are **independent** if and only if

$$\forall a,b \in V: \quad \Pr[\ X=a \ \text{ and } \ Y=b] = \Pr[X=a] \cdot \Pr[Y=b]$$

Example: $U = \{0,1\}^2 = \{00, 01, 10, 11\}$ and $r \xleftarrow{R} U$

Define r.v. X and Y as: $X = \text{lsb}(r)$, $Y = \text{msb}(r)$

$$\Pr[\ X=0 \ \text{ and } \ Y=0\ ] = \Pr[\ r=00\ ] = ¼ = \Pr[X=0] \cdot \Pr[Y=0]$$

Review: XOR

XOR of two strings in $\{0,1\}^n$ is their bit-wise addition mod 2

X	Y	X $\oplus$ Y
0	0	0
0	1	1
1	0	1
1	1	0

$$\begin{array}{r} 0\ 1\ 1\ 0\ 1\ 1\ 1 \\ \oplus \quad 1\ 0\ 1\ 1\ 0\ 1\ 0 \\ \hline 1\ 1\ 0\ 1\ 1\ 0\ 1 \end{array}$$

Independence

- Uniform distribution over $\mathcal{U} = \{0,1\}^2$
- $\mathcal{U}$ = {0,1}^{2} = {00, 01, 10, 11} and r $\xleftarrow{R}$ U
  - X = lsb(r), Y = msb(r), $\mathrm{Z} \coloneqq X + Y$, $W \coloneqq X \oplus Y$

- X, Y independent

- Are X, Z independent?

- Are X, W independent?

An important property of XOR

<u>**Thm**</u>: If Y is a random variable over $\{0,1\}^n$, X is an independent uniform variable on $\{0,1\}^n$

Then $Z := Y \oplus X$ is uniform var. on $\{0,1\}^n$

Proof: (for $n=1$)

$\Pr[\ Z=0\ ] =$

Conditional probability

- For two events A and B, conditional probability is:

$$\Pr[A|B] = \frac{\Pr[A \cap \ B]}{\Pr[B]}$$

- For two random variables X, Y and outcomes x, y we define the conditional probability:

$$\Pr[X = x|Y = y] = \frac{\Pr[X{=}x, Y{=}y]}{\Pr[Y{=}y]}$$

- If A and B are independent

$$\Pr[A|B] = \frac{\Pr[A \cap \ B]}{\Pr[B]} = \frac{\Pr[A]\Pr[B]}{\Pr[B]} = \Pr[\mathrm{A}]$$

Bayes Theorem

- For two events A and B:

$$\Pr[A|B] = \frac{\Pr[B|A]\Pr[A]}{\Pr[B]}$$

- For two random variables X, Y and outcomes x, y

$$\Pr[X = x|Y = y] = \frac{\Pr[Y = y|X = x]\Pr[X = x]}{\Pr[Y = y]}$$

- Easy to infer from definition

$$\Pr[A|B] = \frac{\Pr[A \cap \ B]}{\Pr[B]} = \frac{\Pr[B|A]\Pr[A]}{\Pr[B]}$$

Conditional probability example

- Shift cipher: $\mathcal{K}$ = {0,…,25}, Pr[K = k]=1/26
- Assume that distribution of message is

$$\Pr[M = a] = 0.7; \Pr[M = z] = 0.3$$

- What is the probability that ciphertext is b?
- Solution: $M = a, K = 1$ or $M = z, K = 2$

$$\Pr[M = a, K = 1] = \Pr[M = a] \Pr[k = 1] = 0.7 * \frac{1}{26}$$

$$\Pr[M = z, K = 2] = \Pr[M = z] \Pr[k = 2] = 0.3 * \frac{1}{26}$$

$$\Pr[C = b] = 0.3 * \frac{1}{26} + 0.7 * \frac{1}{26} = \frac{1}{26}$$

Conditional probability example

- Shift cipher: $\mathcal{K}$ = {0,…,25}, Pr[K = k]=1/26
- Assume that distribution of message is

$$\Pr[M = a] = 0.7; \Pr[M = z] = 0.3$$

- What is the probability that message is “a” given that ciphertext is “b”?
- Solution:

$$\Pr[M = a | C = b] = \frac{\Pr[C = b | M = a]\Pr[M = a]}{\Pr[C = b]}$$

$$= \frac{\Pr[K = 1]\Pr[M = a]}{\Pr[C = b]} = \frac{\frac{1}{26} * 0.7}{\frac{1}{26}} = 0.7$$

Defining security of encryption

Encryption setting

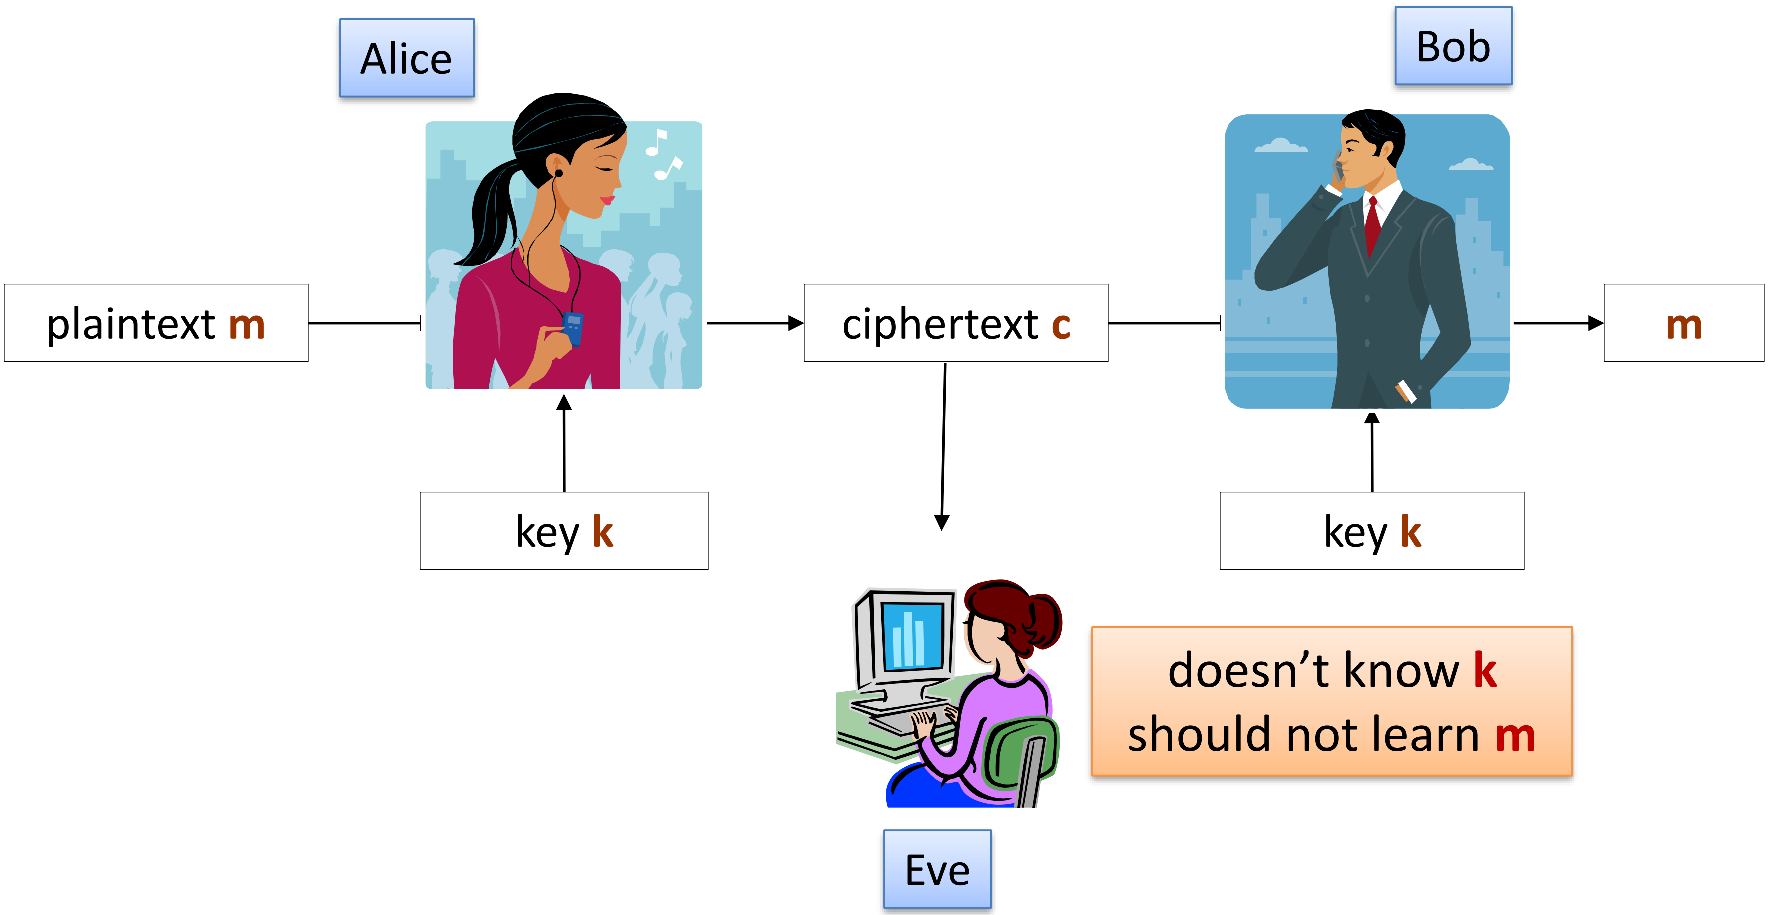

Adversarial capability

- Ciphertext-only attack
  - Adversary observes ciphertext(s)
  - Infer information about plaintext
- Known-plaintext attack
  - Adversary knows one pair of plaintext/ciphertext
  - Learn plaintext information on other ciphertext
- Chosen-plaintext attack
  - Adversary can obtain plaintext/ciphertext pairs of his choice
- Chosen-ciphertext attack
  - Adversary can decrypt ciphertexts of its choice
  - Learn plaintext information on other ciphertext

Defining "security of an encryption scheme" is not trivial.

consider the following experiment

(**m** – a message)

1. the key **K** is chosen uniformly at random
2. **$C := Enc_K(m)$** is given to the adversary

how to define security

?

Idea 1

(m – a message)

1. the key K is chosen uniformly at random
2. $C := Enc_K(m)$ is given to the adversary

An idea

"The adversary should not be able to learn K."

A problem

the encryption scheme that "doesn't encrypt":

$$Enc_K(m) = m$$

satisfies this definition!

Idea 2

(**m** – a message)

1. the key **K** is chosen uniformly at random
2. **$C := Enc_K(m)$** is given to the adversary

An idea

"The adversary should not be able to learn **m**."

A problem

What if the adversary can compute, e.g., the first half of **m**?

m_1	…	$m_{\vert m\vert/2}$	?	…	?

Idea 3

(**m** – a message)

1. the key **K** is chosen uniformly at randomly
2. **C := Enc$_K$(m)** is given to the adversary

An idea

"The adversary should not learn any information about **m**."

Sounds great! But what does it actually mean?
How to formalize it?

Example

m

Eve knows that

$$\mathbf{m} := \begin{cases} \text{“}\textit{I love you}\text{”} & \text{with prob. } \mathbf{0.1} \\ \text{“}\textit{I don’t love you}\text{”} & \text{with prob. } \mathbf{0.7} \\ \text{“}\textit{I hate you}\text{”} & \text{with prob. } \mathbf{0.2} \end{cases}$$

---

m

$c := \mathrm{Enc}_K(m)$

Eve **still** knows that

$$\mathbf{m} := \begin{cases} \text{“}\textit{I love you}\text{”} & \text{with prob. } \mathbf{0.1} \\ \text{“}\textit{I don’t love you}\text{”} & \text{with prob. } \mathbf{0.7} \\ \text{“}\textit{I hate you}\text{”} & \text{with prob. } \mathbf{0.2} \end{cases}$$

Intuitively

"The adversary should not learn any information about **m**."

Consider random variables:

M some distribution variable over $\mathcal{M}$

K uniformly random variable over $\mathcal{K}$

C = Enc(K, M) random variable over $\mathcal{C}$

"The adversary should not learn any information about **m**."

An encryption scheme is **perfectly secret** if

for every distribution of **M**

and every $\mathbf{m} \in \mathcal{M}$ and $\mathbf{c} \in \mathcal{C}$

such that
$P[C = c] > 0$

$$\Pr[\, M = m \,] = \Pr[\, M = m \mid C = c \,]$$

Ciphertext-only attack

Equivalently:

For all **m**, **c**: **Pr[M = m] = Pr[M = m | C = c]**

$\Updownarrow$

M and **C=Enc(K,M)** are independent

$\Updownarrow$

For every **m** , **m'** , **c** we have:
Pr[Enc(K, m) = c] = Pr[Enc(K, m') = c]

One-time pad

A perfectly secret scheme: one-time pad

ℓ – a parameter

$\mathcal{K} = \mathcal{M} = \{0,1\}^{\ell}$

component-wise **xor**

Gilbert Vernam

(1890 –1960)

Vernam's cipher:

$$\mathbf{Enc}_k(m) = k \oplus m$$
$$\mathbf{Dec}_k(c) = k \oplus c$$

Correctness:

$$\mathbf{Dec}_k(\mathbf{Enc}_k(m)) = k \oplus (k \oplus m)$$
$$m$$

Perfect secrecy of the one-time pad

- **Theorem:** The one-time pad satisfies perfect secrecy.
- **Proof:**

Why the one-time pad is not practical?

1. **The key is as long as the message.**
2. **The key cannot be reused.**
3. **Alice and Bob must share a new key every time they communicate**

All three are necessary for perfect secrecy!

This is because:

$$\begin{aligned} Enc_k(m_0) \text{ xor } Enc_k(m_1) &= (k \text{ xor } m_0) \text{ xor } (k \text{ xor } m_1) \\ &= m_0 \text{ xor } m_1 \end{aligned}$$

Key takeaways

- Defining security for encryption is difficult
- Perfect secrecy is one of the first rigorous notion of security
- One-time pad is optimal
  - But many practical drawbacks
  - Still has been used in critical military applications
- Modern cryptography relies on computational assumptions
  - E.g., it is computationally hard to factor large numbers

Acknowledgement

Some of the slides and slide contents are taken from
http://www.crypto.edu.pl/Dziembowski/teaching
and fall under the following:

We have also used materials from Prof. Dan Boneh online cryptography course at Stanford University:

http://crypto.stanford.edu/~dabo/courses/OnlineCrypto/