Instructor: David Wu (dwu4@cs.utexas.edu)
TA: Eli Bradley

Overarching goal of cryptography: securing communication over untrusted networks

Alice → Bob
↓

third party should not be able to
1) eavesdrop of communication (confidentiality)
2) tamper with the communication (integrity)

Today: secure communication on web (https://...)
TLS protocol (transport layer security)
two components: handshake (key exchange)
record layer (confidentiality + integrity)
protecting data at rest: disk encryption

Most of this course: study mechanics for protecting confidentiality + data
- Encryption schemes for confidentiality
- Signature schemes for message integrity
- Key exchange for setting up shared secrets

] "classical" cryptography

Exciting time in cryptography: moving to post-quantum cryptography
↳ This course is also being updated to discuss post-quantum cryptography early (part of unit on public-key encryption)

Logistics and administrivia:
- Course website: https://www.cs.utexas.edu/~dwu4/courses/fa26
- See Ed Discussion for announcements, notes will be posted to course website (1-2 days after lecture)
- Homework submission via Gradescope (enroll via Canvas)
- Course consists of 5 homework assignments (worth 50%) and two in-class exams (worth 50%)
- Five late days for the semester: use in 24-hour increments, max 72 hours (3 late days) for any single assignment
- This is a class on theoretical foundations — focus will be on formally analyzing security of different schemes
  - Will assume comfort with mathematical proofs as well as familiarity with concepts from algorithms and complexity theory (see course prerequisites)
  - Homework + exams are written assignments (no programming component)

A brief history of cryptography:

Original goal was to protect communication (in times of war)

Basic idea: Alice and Bob have a shared key k

Alice computes $c \leftarrow \text{Encrypt}(k, m)$

(c: ciphertext; k: key; m: message (plaintext))

Bob computes $m \leftarrow \text{Decrypt}(k, c)$ to recover the message

This tuple (Encrypt, Decrypt) is called a cipher

K, M, C are sets (e.g., $K = M = C = \{0,1\}^{128}$)

Definition. A cipher is defined over (K, M, C) where K is a key-space, M is a message space and C is a ciphertext space, and consists of two algorithms (Encrypt, Decrypt):

$\text{Encrypt} : K \times M \to C$

$\text{Decrypt} : K \times C \to M$

} functions should be "efficiently-computable"

theory: runs in probabilistic polynomial time [algorithm can be randomized]

practice: fast on an actual computer (e.g., < 10ms on my laptop)

Correctness: $\forall k \in K, \forall m \in M$:

$$\text{Decrypt}(k, \text{Encrypt}(k, m)) = m$$

"decrypting a ciphertext recovers the original message"

Early ciphers:

- Caesar cipher: "shift by 3"

A	↦ D
B	↦ E
C	↦ F
⋮	
X	↦ A
Y	↦ B
Z	↦ C

Not a cipher! There is no key!

Anyone can decrypt!

↳ Algorithm to encrypt is assumed to be public.

NEVER RELY ON SECURITY BY OBSCURITY!

- Harder to change system than a key
- Less scrutiny for secret algorithms

- Caesar cipher++: "shift by k" ($k = 13$: ROT-13)

k is the key

↳ Still totally broken since there are only 26 possible keys (simply via brute force guessing)

- Substitution cipher: the key defines a permutation of the alphabet (i.e., substitution)

A	↦ C
B	↦ X
C	↦ J
⋮	
Z	↦ T

ABC ↦ CXJ

← substitution table is the key

How many keys? For English alphabet, $26! \approx 2^{88}$ possible keys

↑ very large value, cannot brute force the key

Still broken by frequency analysis
- e is the most frequent character (~12%)
- q is the least frequent character (~0.10%)

Can also look at digram, trigram frequencies

- Vigenere cipher (late 1500s) — "polyalphabetic substitution"

key is short phrase (used to determine substitution table):

m = HELLO

k = CAT

Encrypt(k, m):

```
  HELLO
+ CATCA   <- repeat the key
-------
  KFFPP
```

↑ interpret letters as number between 1 and 26
addition is modulo 26

if we know the key length, can break using frequency analysis
otherwise, can try all possible key lengths $l = 1, 2, \ldots$

↳ general assumption: keys will be much shorter than the message (otherwise if we have a good mechanism to deliver long keys securely, then can use that mechanism to share messages directly)

- Fancier substitution ciphers: Enigma (based on rotor machines)

but... still breakable by frequency analysis

Today: encryption done using computers, lots of different ciphers
- AES (advanced encryption standard, 2000) "block cipher"
- Salsa (2005) / ChaCha (2008) "stream cipher"

not ideal property...

One-time pad [Vigenere cipher where key is as long as the message!]

$K = \{0,1\}^n$ Encrypt(k, m): output $c = k \oplus m$

$M = \{0,1\}^n$ Decrypt(k, c): output $m = k \oplus c$

$C = \{0,1\}^n$

↑ bitwise exclusive OR operation (addition mod 2)

Correctness: Take any $k \in \{0,1\}^n$, $m \in \{0,1\}^n$:

$$\text{Decrypt}(k, \text{Encrypt}(k, m)) = k \oplus (k \oplus m) = (k \oplus k) \oplus m = m \quad (\text{since } k \oplus k = 0^n)$$

Is this secure? How do we define security?

- Given a ciphertext, cannot recover the key?
  - NOT GOOD! Says nothing about hiding message. Encrypt$(k, m) = m$ would be secure under this definition, but this scheme is totally insecure intuitively!
- Given a ciphertext, cannot recover the message.
  - NOT GOOD! Can leak part of the message. Encrypt$(k, (m_0, m_1)) = (m_0, m_1 \oplus k)$. This encryption might be considered secure but leaks half the message. [Imagine if message was "username: alice || password: 123456"
    - ↳ this might be the string that is leaked!
- Given a ciphertext, cannot recover any bit of the message.
  - NOT GOOD! Can still learn parity of the bits (or every pair of bits), etc. Information still leaked...
- Given a ciphertext, learn nothing about the message.
  - GOOD! But how to define this?

Coming up with good definitions is difficult! Definitions have to rule out all adversarial behavior (i.e., capture broad enough class of attacks)

↳ Big part of crypto is getting the definitions right. Pre-1970s: cryptography has relied on intuition, but intuition is often wrong! Just because I cannot break it does not mean someone else cannot...

How do we capture "learning nothing about the message"?

If the key is random, then ciphertext should not give information about the message.

Definition. A cipher (Encrypt, Decrypt) satisfies perfect secrecy if for all messages $m_0, m_1 \in M$, and all ciphertexts $c \in C$:

$$\Pr[k \xleftarrow{R} K : \text{Encrypt}(k, m_0) = c] = \Pr[k \xleftarrow{R} K : \text{Encrypt}(k, m_1) = c]$$

probability that encryption of m_0 is c, where the probability is taken over the random choice of the key k

Perfect secrecy says that given a ciphertext, any two messages are equally likely.

⇒ Cannot infer anything about underlying message given only the ciphertext (i.e., "ciphertext-only" attack)

Theorem. The one-time pad satisfies perfect secrecy.

Proof. Take any message $m \in \{0,1\}^n$ and ciphertext $c \in \{0,1\}^n$. Then,

$$\begin{aligned}\Pr[k \xleftarrow{R} \{0,1\}^n : \text{Encrypt}(k,m) = c] &= \Pr[k \xleftarrow{R} \{0,1\}^n : k \oplus m = c] \\ &= \Pr[k \xleftarrow{R} \{0,1\}^n : k = m \oplus c] \\ &= \frac{1}{2^n}\end{aligned}$$

This holds for all messages m and ciphertexts c, so one-time pad satisfies perfect secrecy.