MIT 6.875/6.5620/18.425

Foundations of Cryptography
Lecture 1

Course website: ***https://mit6875.github.io/***

Course Staff

Instructor

Vinod Vaikuntanathan (vinodv@mit)

TAs

Surya Mathialagan(smathi@mit)

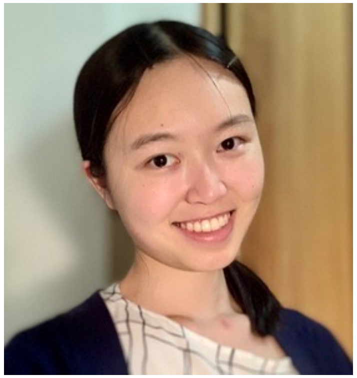

Tina Zhang (tinaz@mit)

Course Website

https://mit6875.github.io/

MIT 6.5620/6.875/18.425 (Fall 2022)
Foundations of Cryptography

Course Description

The field of cryptography gives us a technical language to *define* important real-world problems such as security, privacy and integrity, a mathematical toolkit to *construct* mechanisms such as encryption, digital signatures, zero-knowledge proofs, homomorphic encryption and secure multiparty computation, and a complexity-theoretic framework to *prove* security using reductions that together help us *enforce the rules of the road* in digital interactions.

The last few years have witnessed dramatic developments in the foundations of cryptography, as well as its applications to real-world privacy and security problems. For example, cryptography is abuzz with solutions to long-standing open problems such as fully homomorphic encryption and software obfuscation that use an abundance of data for public good without compromising security.

The course will explore the rich theory of cryptography all the way from the basics to the recent developments.

Prerequisites: This is an introductory, but fast-paced, graduate course, intended for beginning graduate students and upper level undergraduates in CS and Math. We will assume fluency in algorithms (equivalent to 6.046), complexity theory (equivalent to 6.045) and discrete probability (equivalent to 6.042). Mathematical maturity and an ease with writing mathematical proofs will be assumed starting from the first lecture.

Course Information

INSTRUCTOR	Vinod Vaikuntanathan Email: vinodv at csail dot mit dot edu
LOCATION AND TIME	**Monday and Wednesday 1:00-2:30pm** in **1-190**
TAs	**Surya Mathialagan** Email: smathi at mit dot edu Office hours: Monday 3-4pm in TBD, Tuesday 10-11am in 24-323. **Tina Zhang** Email: tinaz at mit dot edu Office hours: Tuesday 5-6pm in TBD, Thursday 4:15-5:15pm in TBD.

Crypto ≠ Cryptocurrencies

6.875 is ***not*** about

Blockchains/
Cryptocurrencies

"Trustworthy"
machine learning

6.875 *is* about foundations:

- Digital Signatures
- Zero-knowledge Proofs
- Public-key Encryption
- Homomorphic Encryption
- Threshold Cryptography
- Pseudorandomness

The Intellectual Origins

Claude E. Shannon

"Communication Theory of Secrecy Systems" (1945)

preceded

"A Mathematical Theory of Communication" (1948)

which founded ***Information Theory***

Cryptanalysis of the Enigma Machine (~1938-39)

"On Computable Numbers, with an Application to the Entscheidungsproblem" (1936)

which gave birth to ***Computer Science***

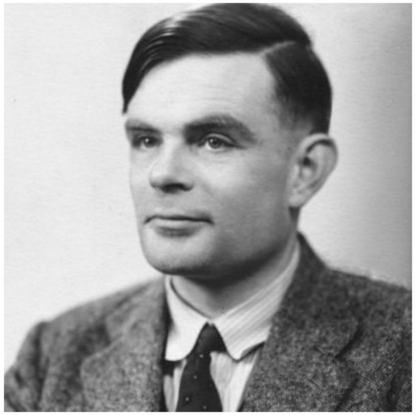

Alan M. Turing

Modern Cryptography:

Practice to Theory and Back

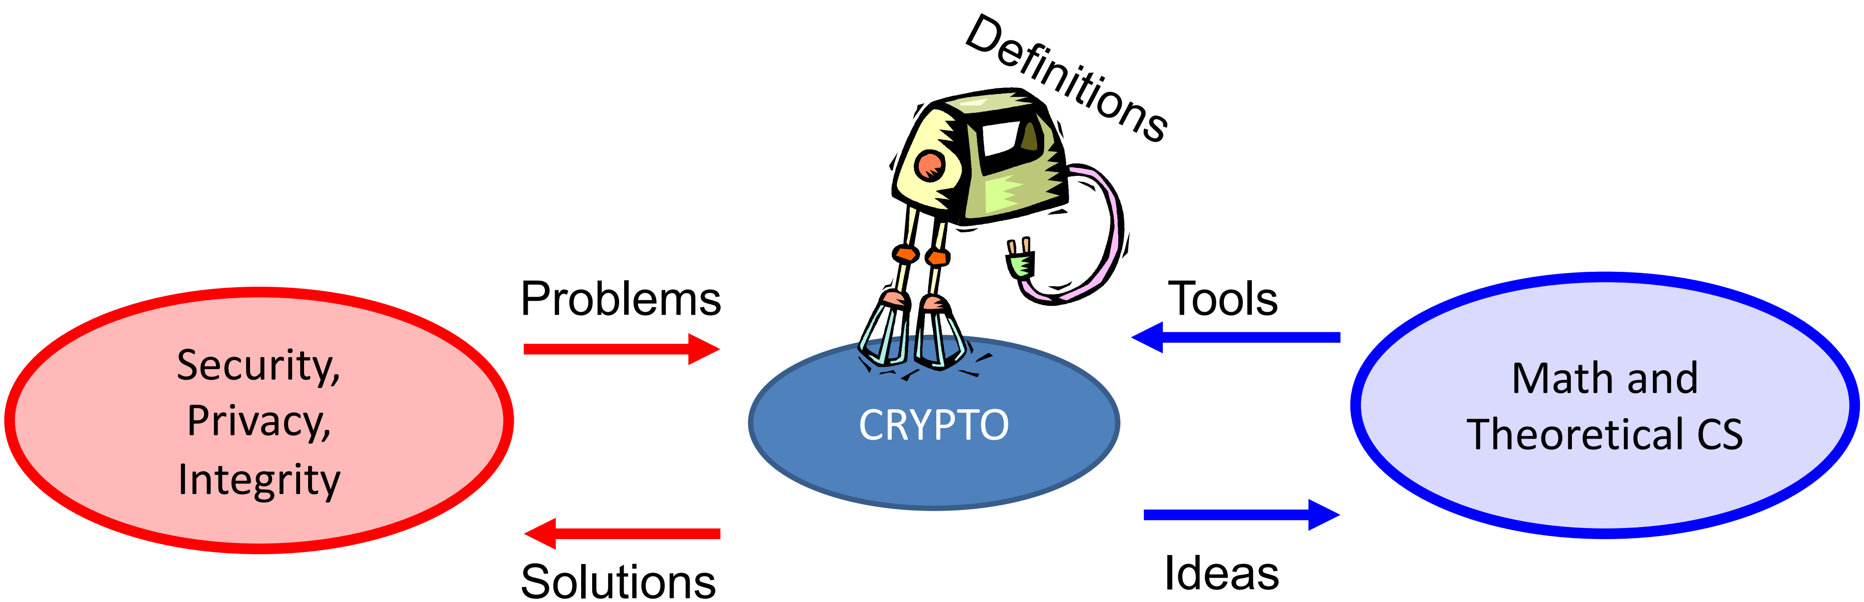

Encryption

Digital Signatures

Pseudorandom Functions

…

Interactive Proofs

Probabilistically checkable Proofs

Locally decodable Codes

…

6.875/6.5620 Themes

1. The Omnipresent, Worst-case, Adversary. 

Central idea. model the adversary: what they know, what they can do, and what their goals are.

Definitions will be our friend.

If you cannot define something, you cannot achieve it.

A key takeaway from 6.875:

Cryptographic (or, adversarial) thinking.

6.875/6.5620 Themes

2. Computational Hardness will be our enabler.
(starting lecture 2)

Central theme: the cryptographic leash. Use computational hardness to “tame” the adversary.

A classical source of hard problems: number theory.

> *“Both Gauss and lesser mathematicians may be justified in rejoicing that there is one such science [number theory] at any rate, whose very remoteness from ordinary human activities should keep it gentle and clean”*
>
> [G. H. Hardy, “A Mathematician’s Apology”]

More recently: geometry, coding theory, combinatorics.

Cryptography is the science of useful hardness.

6.875 Themes

3. Security Proofs via Reductions.

"If there is an (efficient) adversary that breaks scheme A w.r.t. definition D, then there is an (efficient) adversary that factors large numbers."

"Science wins either way"

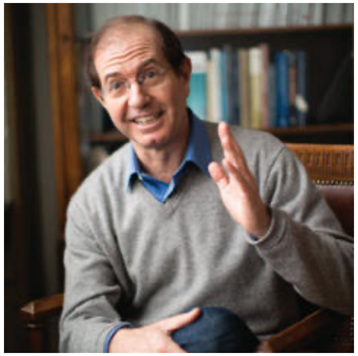

Our reductions will be probabilistic and significantly more involved than the NP-hardness reductions in, say, 6.045.

6.875 Topics

- Pseudorandomness
- Secret-key Encryption and Authentication
- Public-key Encryption and Digital Signatures
- Cryptographic Hashing
- Zero-knowledge Proofs
- Secure Multiparty Computation
- Private Information Retrieval
- Homomorphic Encryption
- Advanced topics:
  Threshold Cryptography, Differential Privacy,…

Administrivia

- **Course website,** the central point of reference.

  https://mit6875.github.io

  Piazza for questions, Gradescope for psets.

  Piazza code: publickey

  Gradescope code: V5DDBG

Administrivia

- **Course website,** the central point of reference.

  https://mit6875.github.io

  Piazza for questions, Gradescope for psets.

- **Homework (95%):** 6 psets, we will count your best 5.

  6.875 is on https://psetpartners.mit.edu

Administrivia

- **Course website,** the central point of reference.

  https://mit6875.github.io

  Piazza for questions, Gradescope for psets.

- **Homework (95%):** 6 psets, we will count your best 5.
- **Class Participation (5%):** Lecture, Recitations, Piazza.
- **Prerequisites:** Algorithms, Probability & Discrete Math, but most of all, “mathematical maturity”.
- **(Optional) special recitations:** 1. probability (this Friday), 2. basic complexity theory, 3. number theory.

Secure Communication

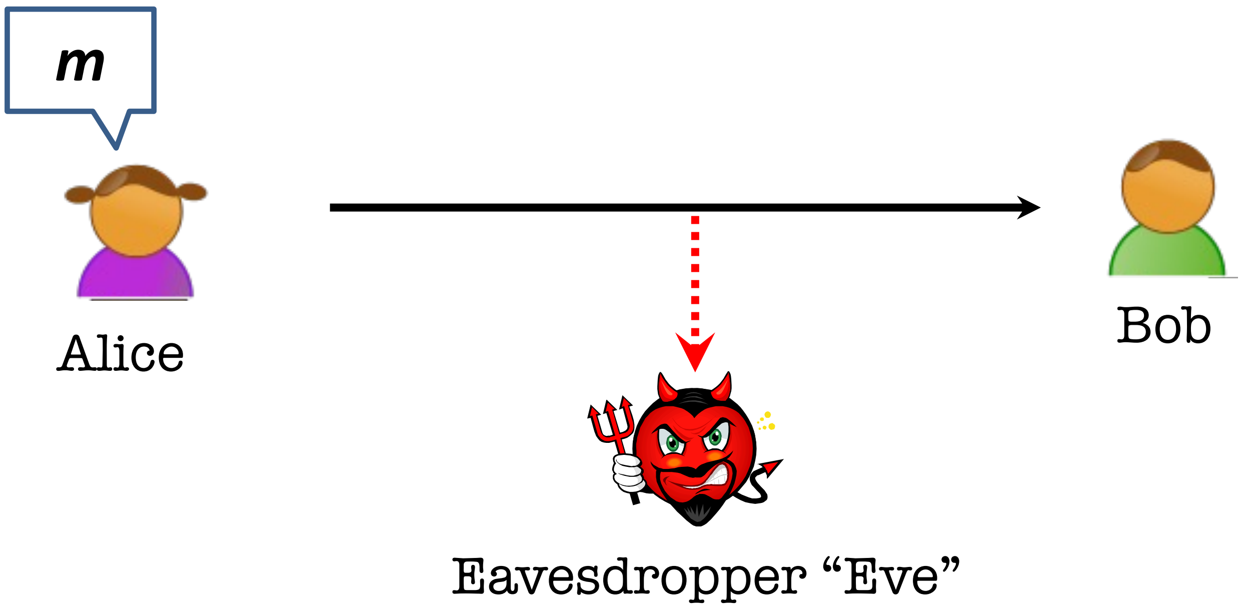

Alice wants to send a message m to Bob without revealing it to Eve.

Secure Communication

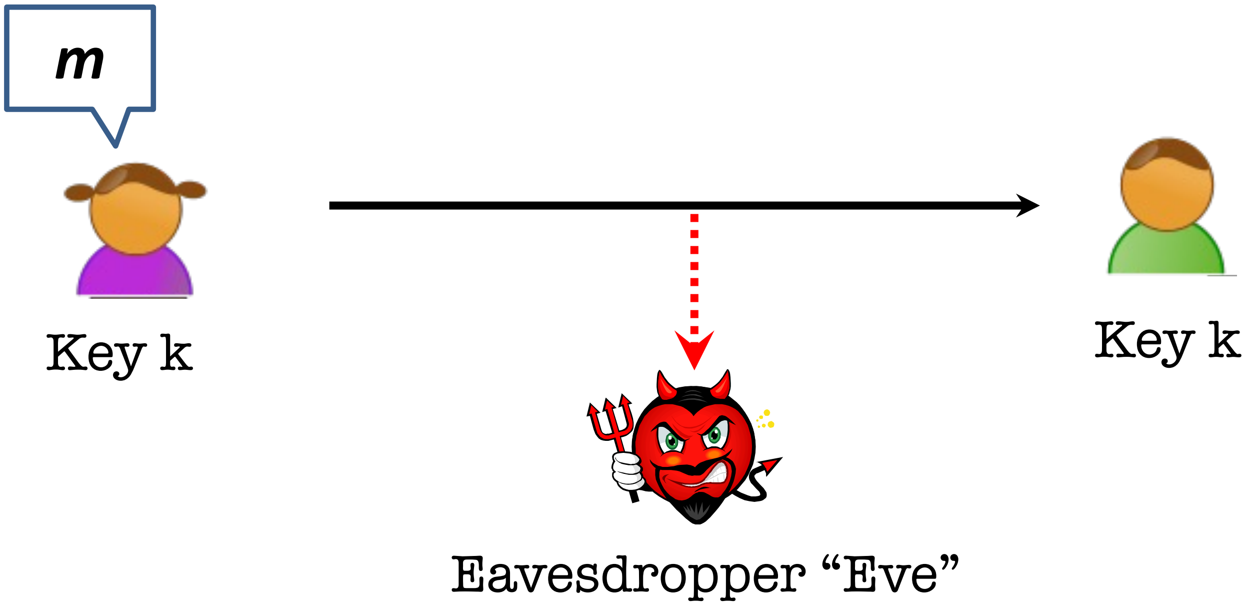

SETUP: Alice and Bob meet beforehand to agree on a secret key k.

Key Notion: Secret-key Encryption

(or Symmetric-key Encryption)

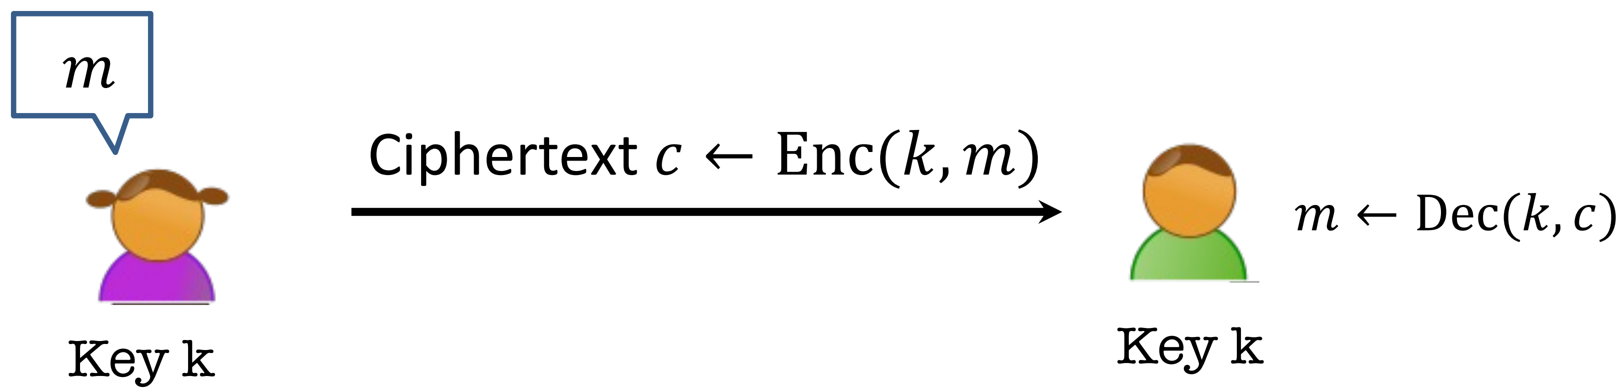

Three (possibly probabilistic) polynomial-time algorithms:

- **Key Generation Algorithm Gen:** $k \leftarrow \mathrm{Gen}(\quad)$
  Has to be probabilistic
- **Encryption Algorithm Enc:** $c \leftarrow \mathrm{Enc}(k, m)$
- **Decryption Algorithm Dec:** $m \leftarrow \mathrm{Dec}(k, c)$

The Worst-case Adversary

- An arbitrary computationally *unbounded* algorithm **EVE**.*

- Knows Alice and Bob's algorithms Gen, Enc and Dec but does not know the key nor their internal randomness.
  (*Kerckhoff's principle or Shannon's maxim*)

- Can see the ciphertexts going through the channel
  (but cannot modify them… we will come to that later)

Security Definition: What is she trying to learn?

Shannon's Perfect Secrecy Definition

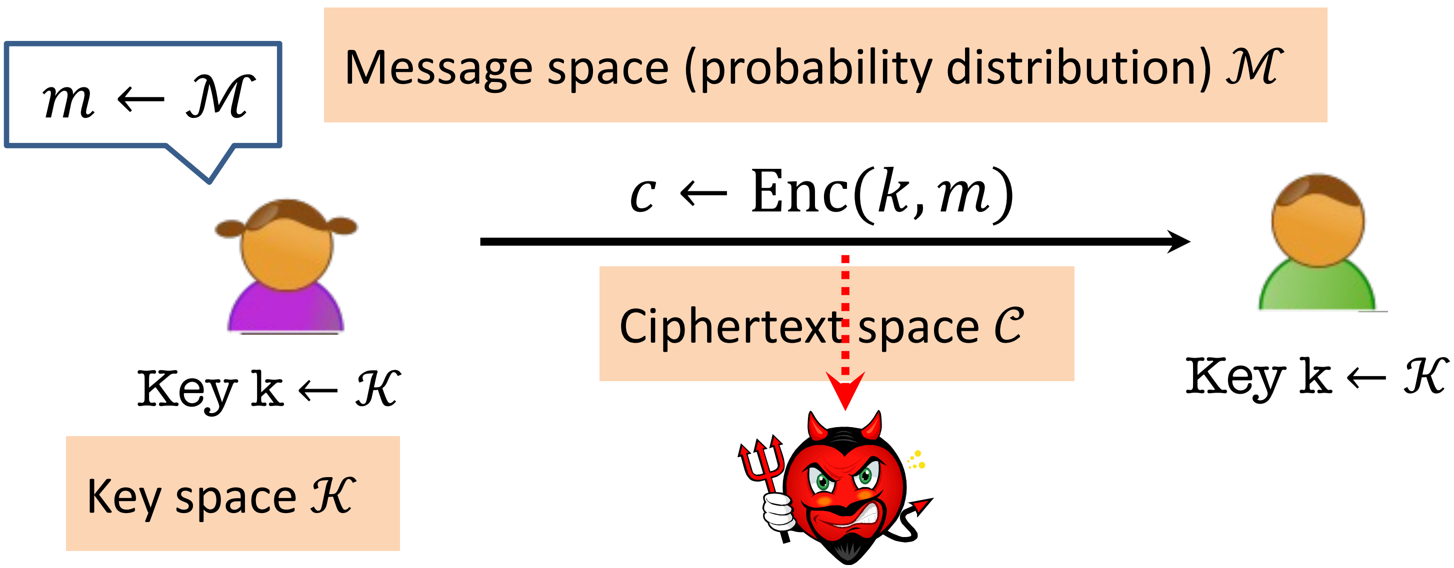

IDEA: A-posteriori = A-priori

Perfect Indistinguishability Definition

Perfect indistinguishability: *a Turing test*

$\forall \mathcal{M}\ \forall m, m' \in \text{Supp}(\mathcal{M}),$

World 0:

$k \leftarrow \mathcal{K}$

$c = E(k, m)$

World 1:

$k \leftarrow \mathcal{K}$

$c' = E(k, m')$

is a **distinguisher** (that gets c and tries to guess which world she's in)

The Two Definitions are Equivalent

THEOREM: An encryption scheme $(\mathrm{Gen}, \mathrm{Enc}, \mathrm{Dec})$ satisfies perfect secrecy IFF it satisfies perfect indistinguishability.

PROOF: Simple use of Bayes' Theorem.

1. Indistinguishability $\Longrightarrow$ Secrecy

WE KNOW (IND): $\forall \mathcal{M}\ \forall m, m' \in \mathrm{Supp}(\mathcal{M}), \forall c \in \mathrm{Supp}(\mathcal{C})$,

$$\Pr[Enc(\mathcal{K}, m) = c] = \Pr[Enc(\mathcal{K}, m') = c] \qquad = \boldsymbol{\alpha}$$

WE WANT (SEC): $\forall \mathcal{M}\ \forall m \in \mathrm{Supp}(\mathcal{M}), \forall c \in \mathrm{Supp}(\mathcal{C})$,

$$\Pr[\mathcal{M} = m | Enc(\mathcal{K}, \mathcal{M}) = c] = \Pr[\mathcal{M} = m]$$

Key Observation: $\forall m\ \Pr[Enc(\mathcal{K}, \mathcal{M}) = c] = \Pr[Enc(\mathcal{K}, m) = c]$.

Proof: *definition of conditional probability*.

$$\Pr[Enc(\mathcal{K}, \mathcal{M}) = c] = \sum \Pr[Enc(\mathcal{K}, \mathcal{M}) = c | \mathcal{M} = m]\ \Pr[\mathcal{M} = m]$$

$$= \sum \Pr[Enc(\mathcal{K}, m) = c] \Pr[\mathcal{M} = m]$$

$$= \alpha \sum \Pr[\mathcal{M} = m] = \alpha.$$

1. Indistinguishability $\Longrightarrow$ Secrecy

WE KNOW (IND): $\forall \mathcal{M}\ \forall m, m' \in \mathrm{Supp}(\mathcal{M}),\ \forall c \in \mathrm{Supp}(\mathcal{C})$,

$$\Pr[Enc(\mathcal{K}, m) = c] = \Pr[Enc(\mathcal{K}, m') = c] \quad = \alpha$$

WE WANT (SEC): $\forall \mathcal{M}\ \forall m \in \mathrm{Supp}(\mathcal{M}),\ \forall c \in \mathrm{Supp}(\mathcal{C})$,

$$\Pr[\mathcal{M} = m | Enc(\mathcal{K}, \mathcal{M}) = c] = \Pr[\mathcal{M} = m]$$

Proof:

$$\Pr[\mathcal{M} = m | Enc(\mathcal{K}, \mathcal{M}) = c] = \frac{\Pr[Enc(\mathcal{K},\mathcal{M})=c|\mathcal{M}=m]\ \Pr[\mathcal{M}=m]}{\Pr[Enc(\mathcal{K},\mathcal{M})=c]} \quad \text{(Bayes)}$$

$$= \frac{\Pr[Enc(\mathcal{K},m)=c]\ \Pr[\mathcal{M}=m]}{\Pr[Enc(\mathcal{K},\mathcal{M})=c]}$$

$$= \frac{\alpha \Pr[\mathcal{M}=m]}{\alpha} = \Pr[\mathcal{M} = m] \quad \text{(key obs.)}$$

2. Secrecy $\Longrightarrow$ Indistinguishability

WE KNOW (SEC): $\forall \mathcal{M}\ \forall m \in \text{Supp}(\mathcal{M}),\ \forall c \in \text{Supp}(\mathcal{C})$,

$$\Pr[\mathcal{M} = m | Enc(\mathcal{K}, \mathcal{M}) = c] = \Pr[\mathcal{M} = m]$$

WE WANT (IND): $\forall \mathcal{M}\ \forall m, m' \in \text{Supp}(\mathcal{M}),\ \forall c \in \text{Supp}(\mathcal{C})$,

$$\Pr[Enc(\mathcal{K}, m) = c] = \Pr[Enc(\mathcal{K}, m') = c]$$

Proof: $\Pr[Enc(\mathcal{K}, m) = c] = \Pr[Enc(\mathcal{K}, \mathcal{M}) = c | \mathcal{M} = m]$

$$= \frac{\Pr[\mathcal{M}=m|Enc(\mathcal{K},\mathcal{M})=c]\ \Pr[Enc(\mathcal{K},\mathcal{M})=c]}{\Pr[\mathcal{M}=m]} \quad \textbf{(Bayes)}$$

$$= \Pr[Enc(\mathcal{K}, \mathcal{M}) = c] \quad \textbf{(because of SEC)}$$

$$= \Pr[Enc(\mathcal{K}, m') = c] \quad \textbf{(symmetry)}$$

Perfect Secrecy is Achievable

The One-time Pad Construction:

Gen: Choose an n-bit string k at random, i.e. $k \leftarrow \{0,1\}^n$

$Enc(k, m)$, where M is an n-bit message: Output $c = m \oplus k$

$Dec(k, c)$: Output $m = c \oplus k$

$\oplus$: bitwise exclusive OR (or XOR)

$0 \oplus 0 = 1 \oplus 1 = 0$

$0 \oplus 1 = 1 \oplus 0 = 1$

$a \oplus b = a + b \pmod 2$

Perfect Secrecy is Achievable

The One-time Pad Construction:

Gen: Choose an n-bit string k at random, i.e. $k \leftarrow \{0,1\}^n$

$Enc(k, m)$, where M is an n-bit message: Output $c = m \oplus k$

$Dec(k, c)$: Output $m = c \oplus k$

Correctness: $c \oplus k = (m \oplus k) \oplus k = m$.

Perfect Secrecy is Achievable

The One-time Pad Construction:

Gen: Choose an n-bit string k at random, i.e. $k \leftarrow \{0,1\}^n$

$Enc(k, m)$, where M is an n-bit message: Output $c = m \oplus k$

$Dec(k, c)$: Output $m = c \oplus k$

Claim: One-time Pad achieves Perfect Indistinguishability (and therefore perfect secrecy).

Proof: For any $m, c \in \{0,1\}^n$,

$\Pr[\text{Enc}(\mathcal{K}, m) = c]$

Perfect Secrecy is Achievable

The One-time Pad Construction:

Gen: Choose an n-bit string k at random, i.e. $k \leftarrow \{0,1\}^n$

$Enc(k, m)$, where M is an n-bit message: Output $c = m \oplus k$

$Dec(k, c)$: Output $m = c \oplus k$

Claim: One-time Pad achieves Perfect Indistinguishability (and therefore perfect secrecy).

Proof: For any $m, m', c \in \{0,1\}^n$,

$$\Pr[\mathrm{Enc}(\mathcal{K}, m) = c] = \Pr[m \oplus \mathcal{K} = c] = \Pr[\mathcal{K} = c \oplus m] = 1/2^n$$

$$= \Pr[\mathcal{K} = c \oplus m'] = \Pr[m' \oplus \mathcal{K} = c] = \Pr[\mathrm{Enc}(\mathcal{K}, m') = c]$$

Perfect Secrecy is Achievable

The One-time Pad Construction:

Gen: Choose an n-bit string k at random, i.e. $k \leftarrow \{0,1\}^n$

$Enc(k, m)$, where M is an n-bit message: Output $c = m \oplus k$

$Dec(k, c)$: Output $m = c \oplus k$

Claim: One-time Pad achieves Perfect Indistinguishability (and therefore perfect secrecy).

Proof: For any $m, m', c \in \{0,1\}^n$,

So, $\Pr[\mathrm{Enc}(\mathcal{K}, m) = c] = \Pr[\mathrm{Enc}(\mathcal{K}, m') = c]$.

QED.

Reusing a One-time Pad?

_m0_

$c0 = m0 \oplus k$

Key k

Key k

A week later:

_m1_

$\mathrm{c1} = m1 \oplus k$

Super-secure Whisper room

Is this *still* perfectly secret?

Reusing a One-time Pad?

Claim: Two-time Pad does ***not*** achieve Perfect Indistinguishability (and therefore ***not*** perfect secrecy).

Proof: Perfect indistinguishability requires that for all pairs $(m0, m1), (m0', m1'), (c0, c1) \in \{0,1\}^{2n}$:

$$\Pr[\mathrm{Enc}(\mathrm{k}, m0) = c0 \; and \; \mathrm{Enc}(\mathrm{k}, m1) = c1]$$
$$= \Pr[\mathrm{Enc}(\mathrm{k}, m0') = c0 \; and \; \mathrm{Enc}(k, m1') = c1]$$

Reusing a One-time Pad?

Claim: One-time Pad does ***not*** achieve Perfect Indistinguishability (and therefore ***not*** perfect secrecy).

Proof: We want to pick $(m0, m1), (m0', m1'), (c0, c1) \in \{0,1\}^{2n}$ s.t.

$$\Pr[\text{Enc}(\text{k}, m0) = c0 \; and \; \text{Enc}(k, m1) = c1]$$
$$\neq \Pr[\text{Enc}(\text{k}, m0') = c0 \; and \; \text{Enc}(k, m1') = c1]$$

Pick $\boldsymbol{m0 = m1 = m}$, $\boldsymbol{m0' \neq m1'}$ **and** $\boldsymbol{c0 = c1 = c}$.

$$\Pr[\text{Enc}(\text{k}, m0) = c0 \; and \; \text{Enc}(k, m1) = c1]$$
$$= \Pr[\text{Enc}(k, m) = c] = 1/2^n$$

Reusing a One-time Pad?

Claim: One-time Pad does ***not*** achieve Perfect Indistinguishability (and therefore ***not*** perfect secrecy).

Proof: We want to pick $(m0, m1), (m0', m1'), (c0, c1) \in \{0,1\}^{2n}$ s.t.

$$\Pr[\text{Enc}(\text{k}, m0) = c0 \; and \; \text{Enc}(k, m1) = c1]$$
$$\neq \Pr[\text{Enc}(\text{k}, m0') = c0 \; and \; \text{Enc}(k, m1') = c1]$$

Pick $\boldsymbol{m0 = m1 = m}$, $\boldsymbol{m0' \neq m1'}$ and $\boldsymbol{c0 = c1 = c}$.

$$\Pr[\text{Enc}(\text{k}, m0) = c = \text{Enc}(k, m1)] = 1/2^n$$

$$\Pr[\text{Enc}(\text{k}, m0') = \text{c} = \text{Enc}(k, m1')] = 0$$

$\blacksquare$

Perfect Secrecy has its Price

THEOREM: For any perfectly secure encryption scheme,
$$|\mathcal{K}| \geq |\mathcal{M}|$$

PROOF (by picture): Assume for contradiction that $|\mathcal{K}| < |\mathcal{M}|$.

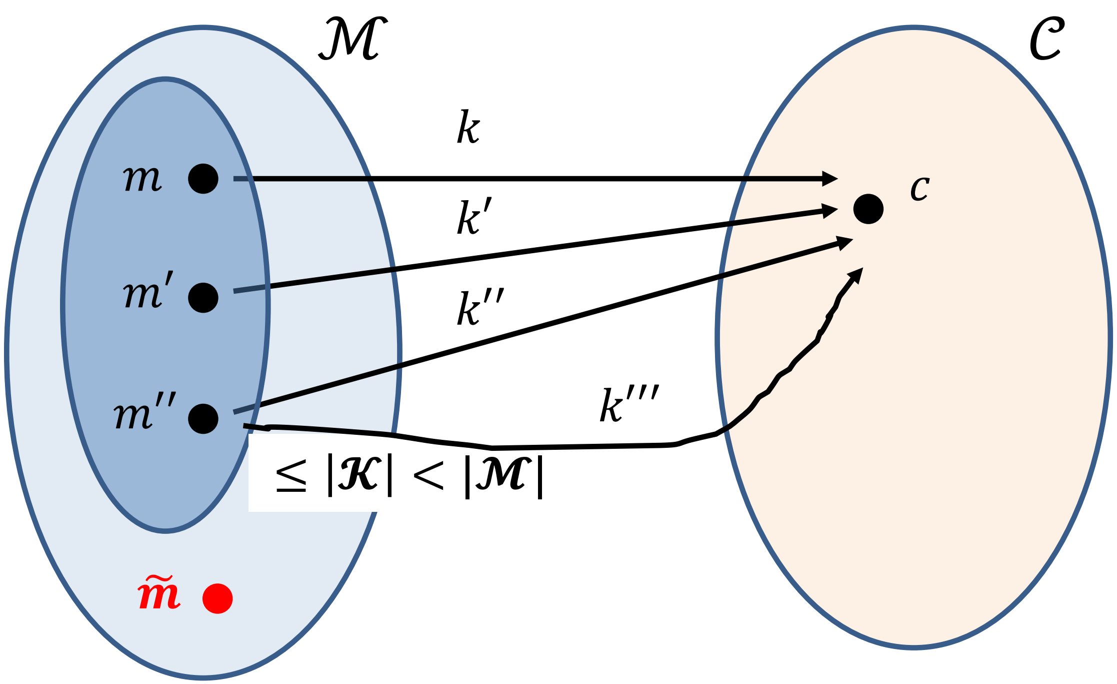

Pick any $c \in \mathcal{C}$

Look at the set of possible msgs ($m = Dec(k, c)$ etc.)

Distinct keys!

Perfect Secrecy has its Price

THEOREM: For any perfectly secure encryption scheme,

$$|\mathcal{K}| \geq |\mathcal{M}|$$

PROOF (by picture): Assume for contradiction that $|\mathcal{K}| < |\mathcal{M}|$.

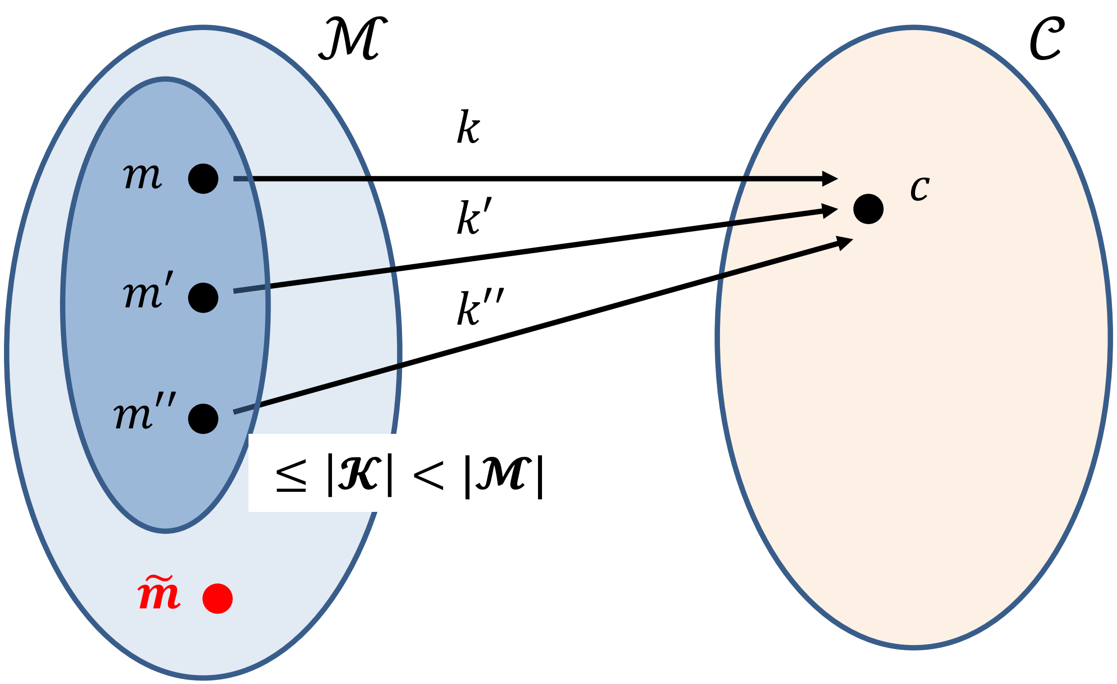

$$\Pr[Enc(\mathcal{K}, m) = c] > 0$$

$$\Pr[Enc(\mathcal{K}, \tilde{m}) = c] = 0$$

***QED*.**

So, what are we to do?

RELAX the definition:

EVE is an arbitrary *computationally bounded* algorithm.

+ number theory/geometry/combinatorics

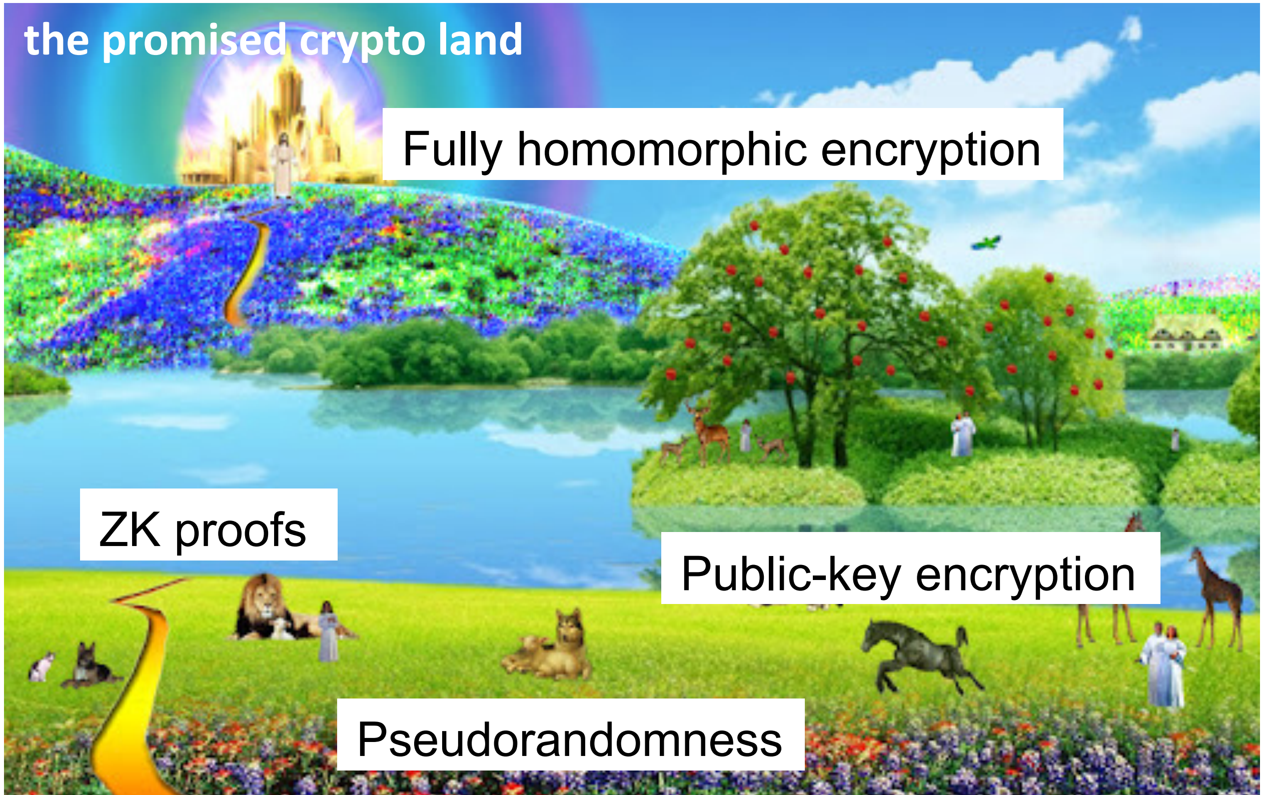

To Summarize…

- **Secure Communication**: a quintessential problem in cryptography.
- We saw two equivalent definitions of security:
  Shannon's perfect indistinguishability and perfect secrecy
- **One-time pad achieves perfect secrecy.**
- **A Serious Limitation:** Any perfectly secure encryption scheme needs keys that are at least as long as the messages.
- **Next Lecture: Overcoming the limitation** with Computationally Bounded Adversaries.