

CS 65500

Advanced Cryptography

Lecture 1: Basics of Indistinguishability

Instructor: Aarushi Goel

Spring 2025


Agenda

- Independence / Perfect Secrecy
- Statistical Indistinguishability
- Computational Indistinguishability

Private Communication



Alice



Bob

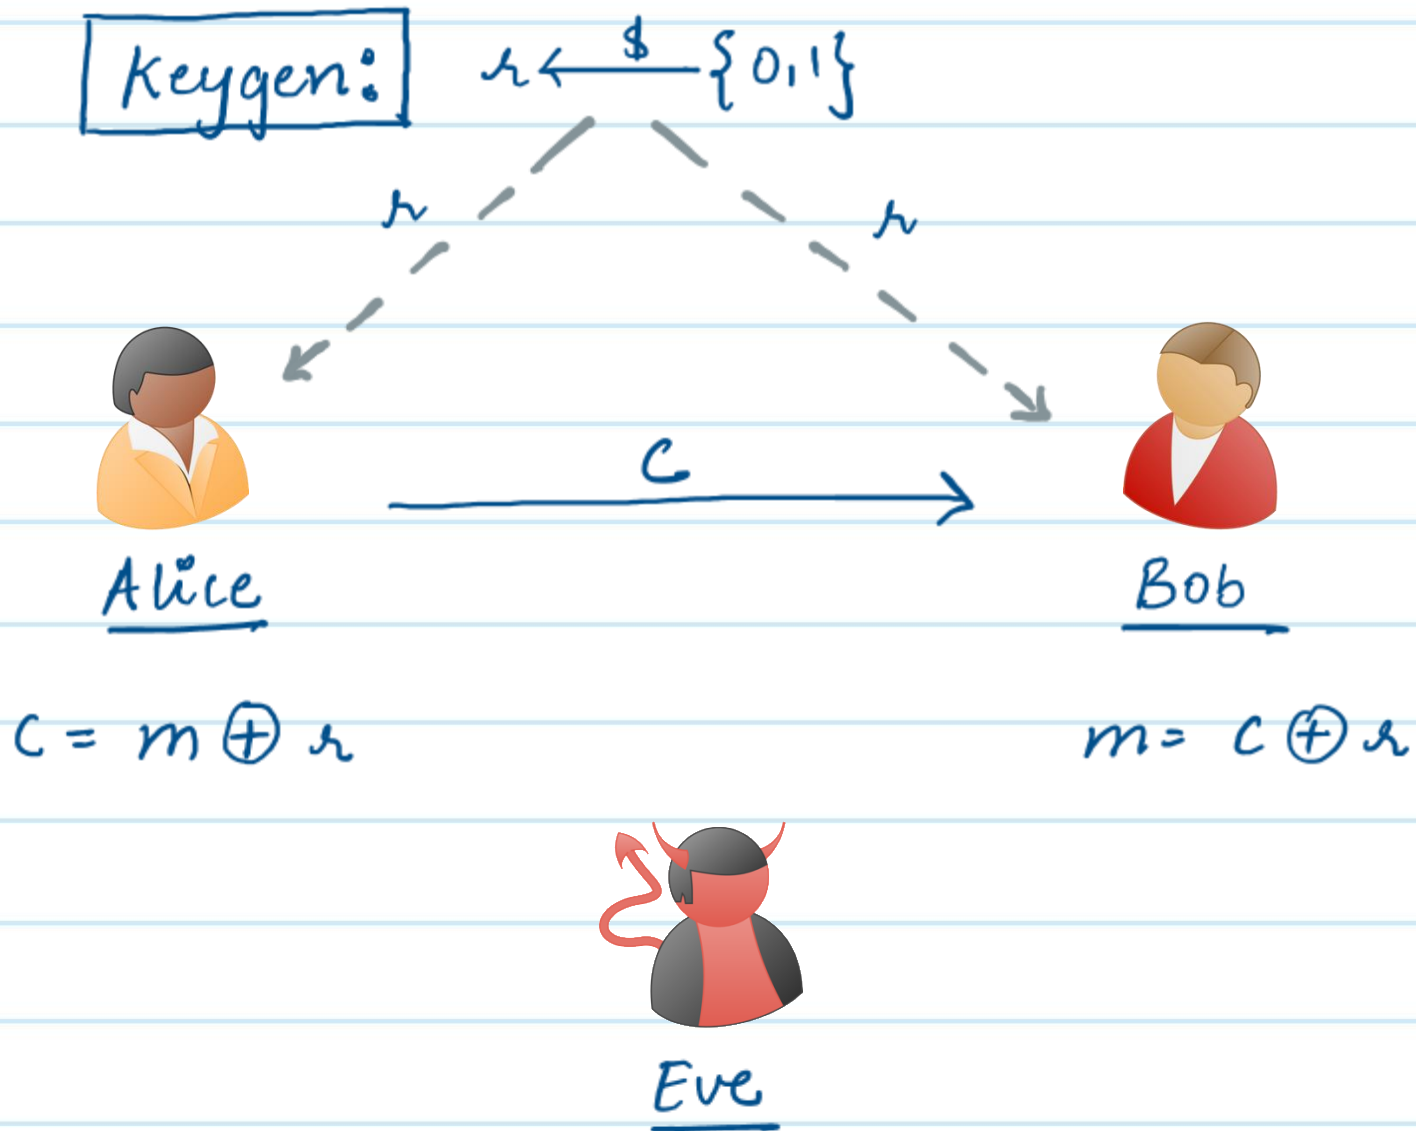


Eve

(Computationally
Unbounded)

How can Alice send m to Bob, while keeping it hidden from an eavesdropper Eve?

One-Time Pad



One-Time Pad

Let $m=0$. What are the possible values of c ?

prob	r	$c = m \oplus r$
$\frac{1}{2}$	0	0
$\frac{1}{2}$	1	1

\Rightarrow Whatever Eve sees
(i.e., c) is independent of m

also called the "view" of
the adversary

Secrecy

→ Is the message m really secret?

→ Eve could have easily guessed m with probability $\frac{1}{2}$.
In fact if they already knew something about m , they can do better.

NOTE: we did not claim that m is random

→ But Eve could have done this without looking at c .
 c did not leak any additional information about m .

Secrecy

Typical goal in cryptography:

PRESERVE SECRECY!

Intuitively speaking, this is what we want:

What Eve learns about m after seeing c ,
is the same as what they already
knew about m .

Formalizing Secrecy

Event 1: What did Eve already know about the message?

- Probability distribution over m
- i.e., $\forall m, \Pr[\text{msg} = m]$

Event 2: What does Eve learn after seeing c ?

- New distribution $\Pr[\text{msg} = m \mid \text{view} = c]$

What do we want for secrecy?

Eve's knowledge in Event 1
=

Eve's knowledge in Event 2

Formalizing Secrecy

$$\Rightarrow \forall m, \forall c, \Pr[\text{msg} = m \mid \text{view} = c] = \Pr[\text{msg} = m]$$

view is independent of msg

$$\Rightarrow \forall m, \forall c, \Pr[\text{view} = c \mid \text{msg} = m] = \Pr[\text{view} = c]$$

for all possible values of msg, the view is distributed identically

$$\Rightarrow \forall c, \forall m_1, m_2, \Pr[\text{view} = c \mid \text{msg} = m_1] = \Pr[\text{view} = c \mid \text{msg} = m_2]$$

Formalizing Secrecy (Summary)

These are equivalent formulations:

$$\forall m, \forall c, \Pr[\text{msg} = m \mid \text{view} = c] = \Pr[\text{msg} = m]$$

$$\forall m, \forall c, \Pr[\text{view} = c \mid \text{msg} = m] = \Pr[\text{view} = c]$$

$$\forall m, \forall c, \Pr[\text{msg} = m, \text{view} = c] = \Pr[\text{msg} = m] \times \Pr[\text{view} = c]$$

$$\forall c, \forall m_1, m_2, \Pr[\text{view} = c \mid \text{msg} = m_1] = \Pr[\text{view} = c \mid \text{msg} = m_2]$$

Formalizing Secrecy

Is $\Pr[\text{msg} = m_1 \mid \text{view} = v] = \Pr[\text{msg} = m_2 \mid \text{view} = v]$?

Why / Why not?

This is only true if msg is uniform.

Relaxing Secrecy Requirement

What if the view is not exactly independent of the message?

Next Best Thing:

view is close to a distribution that is independent of the message

- Statistical Closeness
- Computational Closeness

Statistical Difference

Given two distributions A & B over some sample space, how well can a test T distinguish between them?

- T is given a sample drawn from A or B
- How differently does it behave in the two cases?

$$\Delta(A, B) = \max_T \left| \Pr_{x \leftarrow A} [T(x) = 0] - \Pr_{x \leftarrow B} [T(x) = 0] \right|$$

↓
Statistical
difference
between
 A & B

↓
max over all such
possible tests

Statistical Indistinguishability

→ A and B are statistically indistinguishable from each other if the statistical difference between them is negligible. Examples: 2^{-20} , 2^{-50} , 2^{-100}
we let the user decide which of these they want.

→ Decide how? Using a security parameter k
Security guarantees will be given asymptotically as a function of the security parameter.

→ Given $\{A_k\}$ & $\{B_k\}$, $\Delta(A_k, B_k)$ is a function of k .

we want this to be negligible function in k !

Negligible Functions

- The best distinguishing test T^* should have extremely small probability of success.
- If T^* has extremely small success probability given one sample, the best distinguishing test should also have extremely small success probability given polynomially many samples.
- Functions that decay so quickly (i.e., approach to zero) that they cannot be rescued by any polynomial.

Negligible Functions

Definition: A function $v(\cdot)$ is negligible if for every polynomial $p(\cdot)$, we have

$$\lim_{n \rightarrow \infty} p(n) \cdot v(n) = 0$$

\Rightarrow A negligible function decays faster than all inverse polynomial functions.

Definition: A function $v(k)$ is negligible if $\forall c \geq 0, \exists N$
s.t., $\forall k > N, v(k) \leq \frac{1}{k^c}$

Statistical Indistinguishability

Definition: Distribution ensembles $\{A_k\}$, $\{B_k\}$ are statistically indistinguishable if \exists negligible $\nu(\cdot)$, s.t., $\forall k$, $\Delta(A_k, B_k) \leq \nu(k)$

$\Rightarrow \exists$ negligible $\nu(\cdot)$, s.t. \forall tests T , $\forall k$

$$\left| \Pr_{x \leftarrow A_k} [T_k(x) = 0] - \Pr_{x \leftarrow B_k} [T_k(x) = 0] \right| \leq \nu(k)$$

Ques: is this equivalent to: \forall tests T , \exists negligible $\nu(\cdot)$ s.t. $\forall k$ $\left| \Pr_{x \leftarrow A_k} [T_k(x) = 0] - \Pr_{x \leftarrow B_k} [T_k(x) = 0] \right| \leq \nu(k)$?

Computational Indistinguishability

Definition: Distribution ensembles $\{A_k\}, \{B_k\}$ are computationally indistinguishable if \nexists efficient tests T , \exists negligible $\nu(\cdot)$. s.t. $\forall k$,

$$\left| \Pr_{x \leftarrow A_k} [T_k(x) = 0] - \Pr_{x \leftarrow B_k} [T_k(x) = 0] \right| \leq \nu(k)$$

What is efficient?

Cost of Computation

It can be helpful to think of cost of computation in terms of monetary value. Following costs are approximated using the pricing model of Amazon EC2

<i>clock cycles</i>	<i>approx cost</i>	<i>reference</i>
2^{50}	\$3.50	<i>cup of coffee</i>
2^{55}	\$100	<i>decent tickets to a Portland Trailblazers game</i>
2^{65}	\$130,000	<i>median home price in Oshkosh, WI</i>
2^{75}	\$130 million	<i>budget of one of the Harry Potter movies</i>
2^{85}	\$140 billion	<i>GDP of Hungary</i>
2^{92}	\$20 trillion	<i>GDP of the United States</i>
2^{99}	\$2 quadrillion	<i>all of human economic activity since 300,000 BC⁴</i>
2^{128}	<i>really a lot</i>	<i>a billion human civilizations' worth of effort</i>

Computational Security



John Nash

“It doesn't really matter whether attacks are impossible, only whether attacks are computationally infeasible”

Modern cryptography is based on this principle.

Efficient = Probabilistic Polynomial Time. (PPT)

Computational Security

Non-uniform PPT: A family of randomized programs $\{T_k\}$ (one for each value of the security parameter k), s.t. there is a polynomial $p(\cdot)$ with each T_k running for at most time $p(k)$.

Uniform PPT: where T is a single program that takes k as an additional input.

By default: We will consider non-uniform PPT algorithms / adversaries / tests / distinguishers.

Choosing an Appropriate Security Parameter

Some example references for what extremely small probabilities are equivalent to.

<i>probability</i>	<i>equivalent</i>
--------------------	-------------------

2^{-10}	<i>full house in 5-card poker</i>
-----------	-----------------------------------

2^{-20}	<i>royal flush in 5-card poker</i>
-----------	------------------------------------

2^{-28}	<i>you win this week's Powerball jackpot</i>
-----------	----------------------------------------------

2^{-40}	<i>royal flush in 2 consecutive poker games</i>
-----------	-------------------------------------------------

2^{-60}	<i>the next meteorite that hits Earth lands in this square →</i>
-----------	------------------------------------------------------------------

