CSCI 1510

- ZKP for All NP (continued)
- Non-Interactive Zero-Knowledge Proofs
- Definitions of Secure Multi-Party Computation

Zero-Knowledge Proof (ZKP)

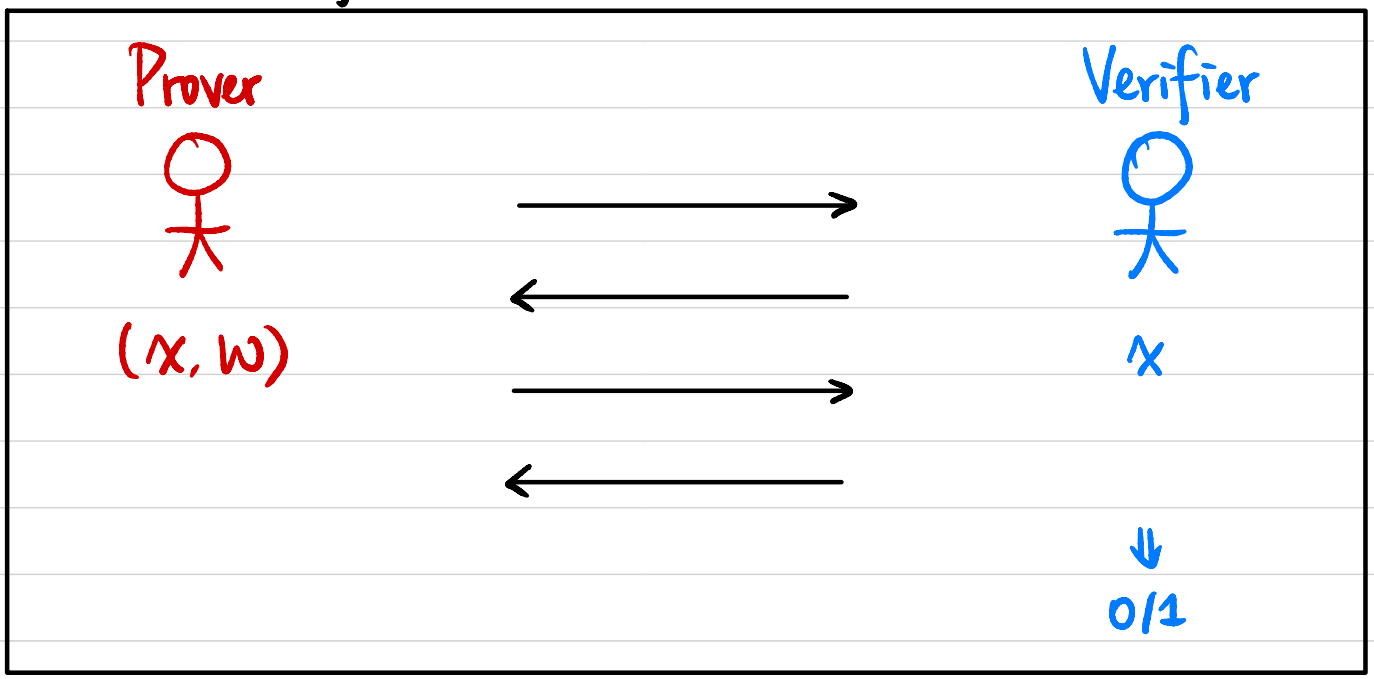

Let (P, V) be a pair of PPT interactive machines. (P, V) is a **zero-knowledge proof system** for a language L with associated relation R_L if

- **Completeness:** $\forall (x, w) \in R_L, \quad Pr[\, P(x,w) \longleftrightarrow V(x) \text{ outputs } 1 \,] = 1.$
- **Soundness:** $\forall x \notin L, \quad \forall \underset{\text{argument}}{(PPT)} P^*, \quad Pr[\, P^*(x) \longleftrightarrow V(x) \text{ outputs } 1 \,] \leq negl(n).$
- **Zero-Knowledge:** $\forall PPT\ V^*, \quad \exists PPT\ S \quad s.t. \quad \forall (x, w) \in R_L,$

$$Output_{V^*}[P(x,w) \longleftrightarrow V^*(x)] \simeq S(x)$$

ZKP for Graph 3-Coloring (All NP)

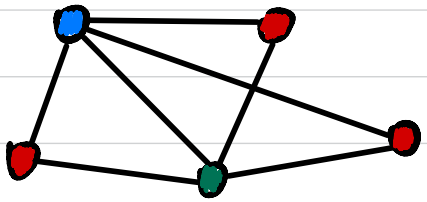

NP language $L = \{ G : G \text{ has 3-coloring} \}$

NP relation $R_L = \{ (G, 3COL) \}$

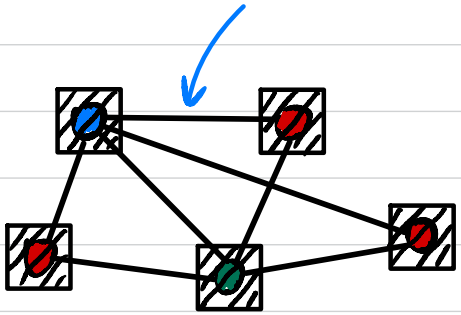

$\pi : \{ \bullet \; \bullet \; \bullet \} \rightarrow \{ \bullet \; \bullet \; \bullet \}$

Commitment Scheme

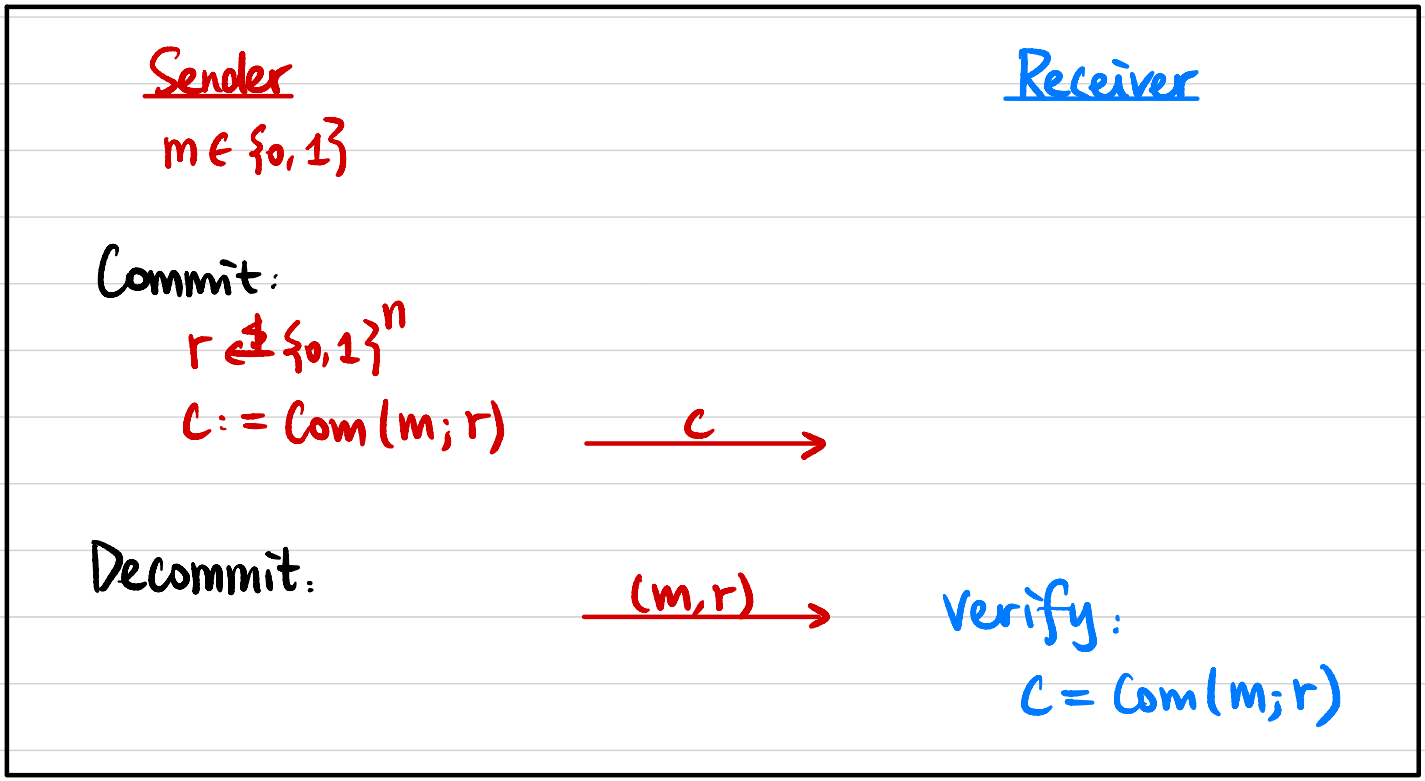

- **Perfectly Binding:** $\forall r, s \in \{0,1\}^n,\ Com(0; r) \neq Com(1; s)$

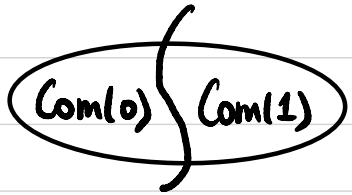

- **Computationally Hiding:** $Com(0; U_n) \overset{c}{\approx} Com(1; U_n)$

ZKP for Graph 3-Coloring

Input: $G = (V, E)$

Witness: $\phi: V \to \{0,1,2\}$

Given a perfectly binding commitment scheme Com.

Soundness?

G & L, by perfect binding of Com,

$\Pr[P^* \text{ not caught}] \le (1 - \frac{1}{|E|})^{n \cdot |E|}$

$\approx e^{-n}$

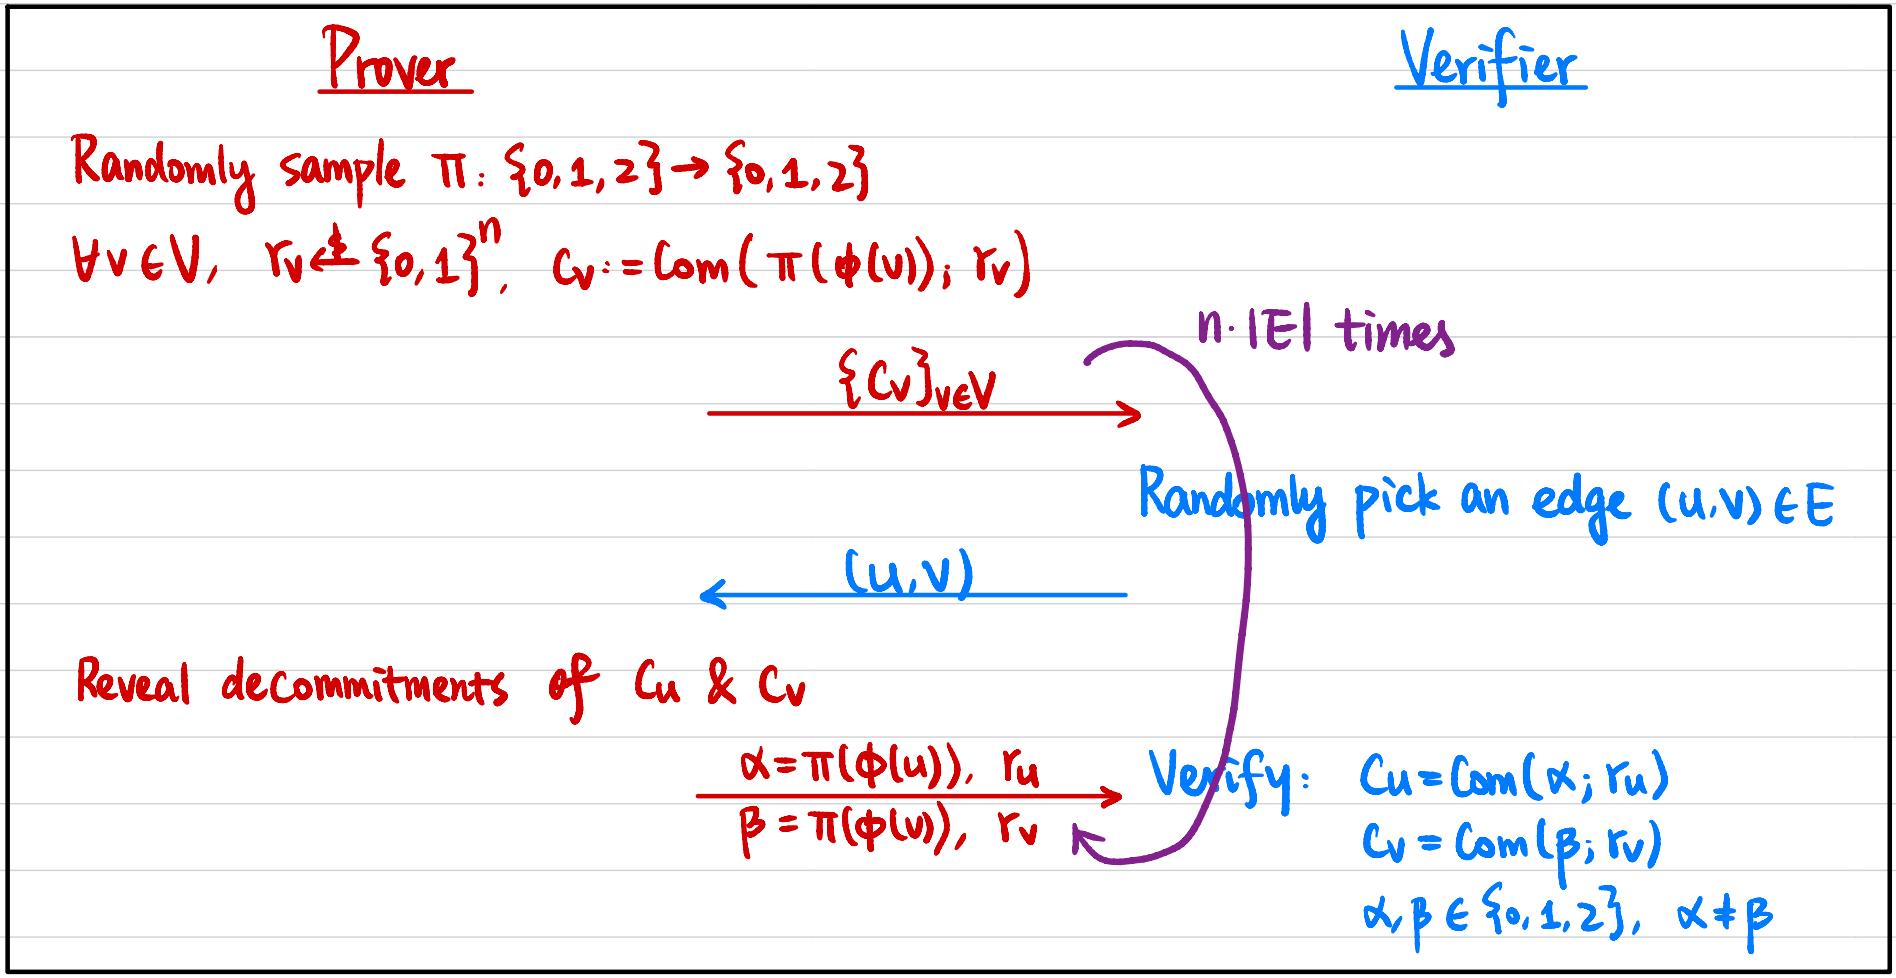

Completeness?

Zero-Knowledge?

$\forall$ PPT V^*, $\exists$ PPT S s.t. $\forall (x, w) \in R_L$,

$$\text{Output}_{V^*}[P(x,w) \longleftrightarrow V^*(x)] \overset{c}{\approx} S(x)$$

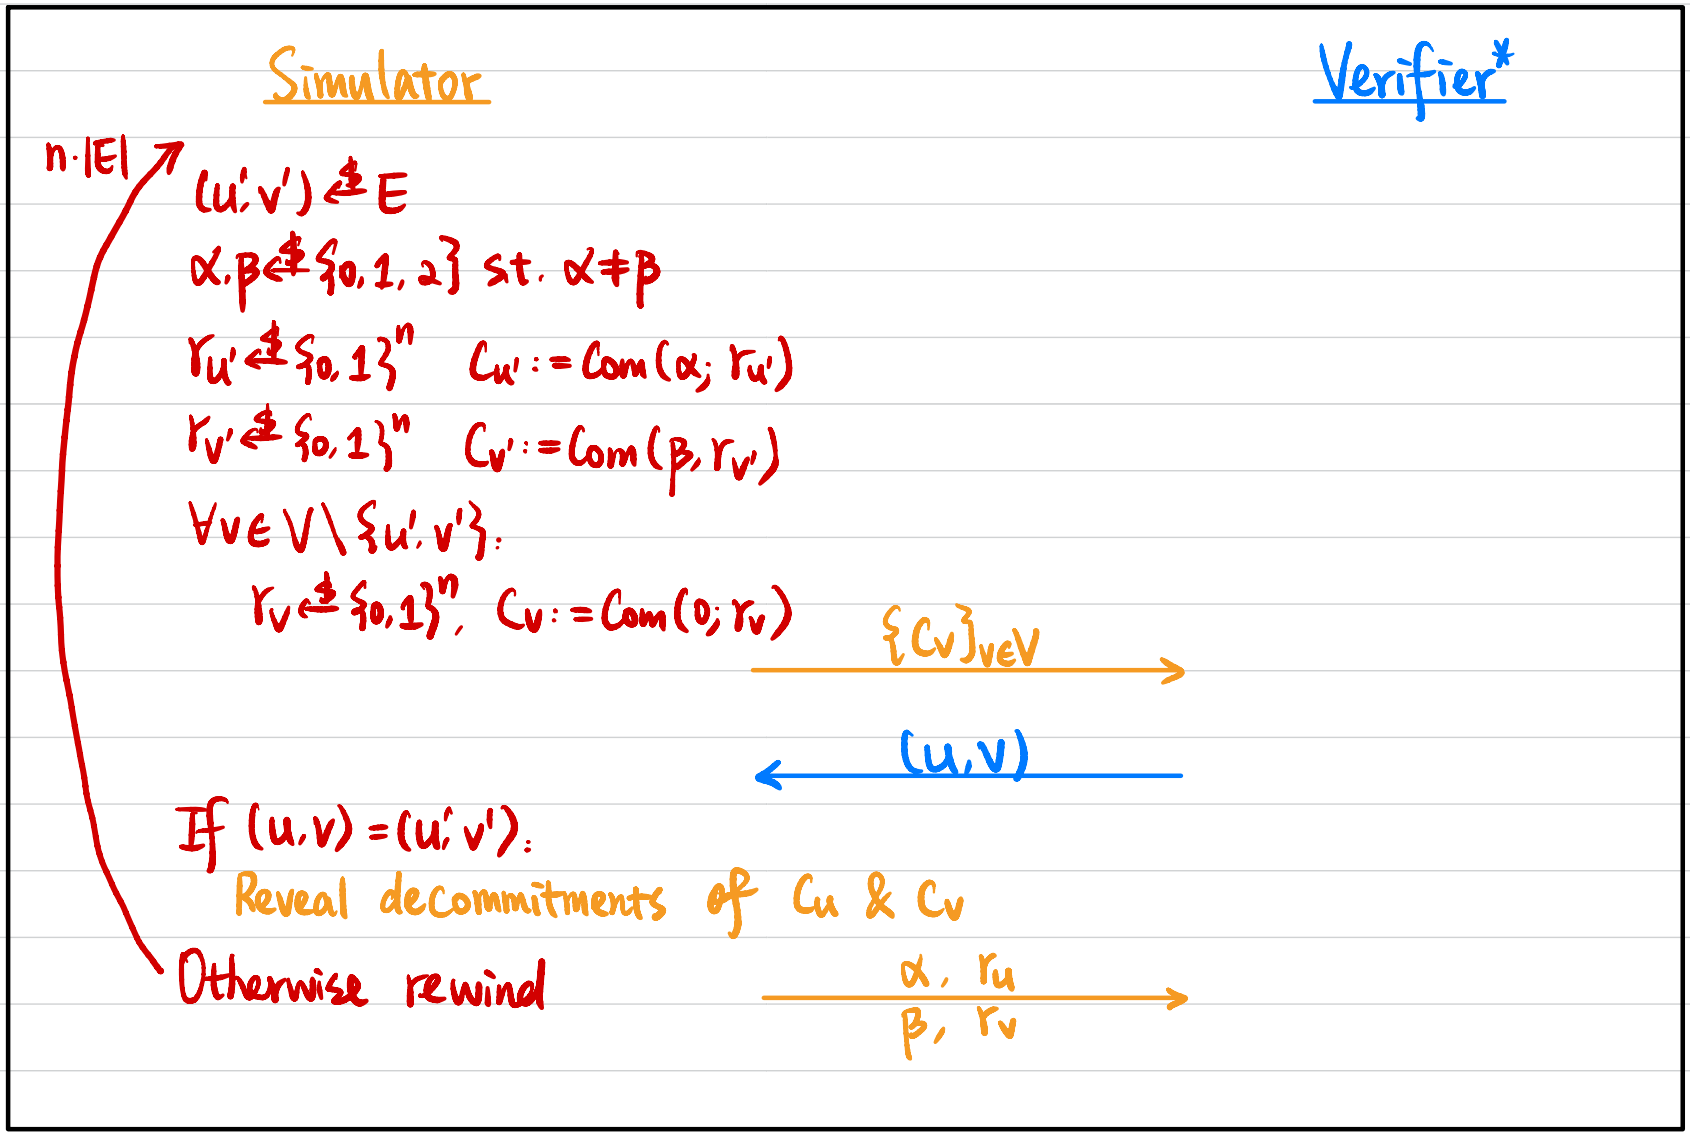

Prover	$\mathcal{H}_1$	Verifier*

$(u', v') \xleftarrow{\$} E$

$\alpha, \beta \xleftarrow{\$} \{0,1,2\}$ s.t. $\alpha \neq \beta$

Construct $\pi : \{0,1,2\} \to \{0,1,2\}$ s.t.

$\pi(\phi(u')) = \alpha \;\wedge\; \pi(\phi(v')) = \beta$

$\forall v \in V,\ r_v \xleftarrow{\$} \{0,1\}^n,\ C_v := \mathsf{Com}(\pi(\phi(v)); r_v)$

$\{C_v\}_{v \in V}$ →

← (u, v)

Reveal decommitments of C_u & C_v

$\alpha = \pi(\phi(u)),\ r_u$

$\beta = \pi(\phi(v)),\ r_v$ →

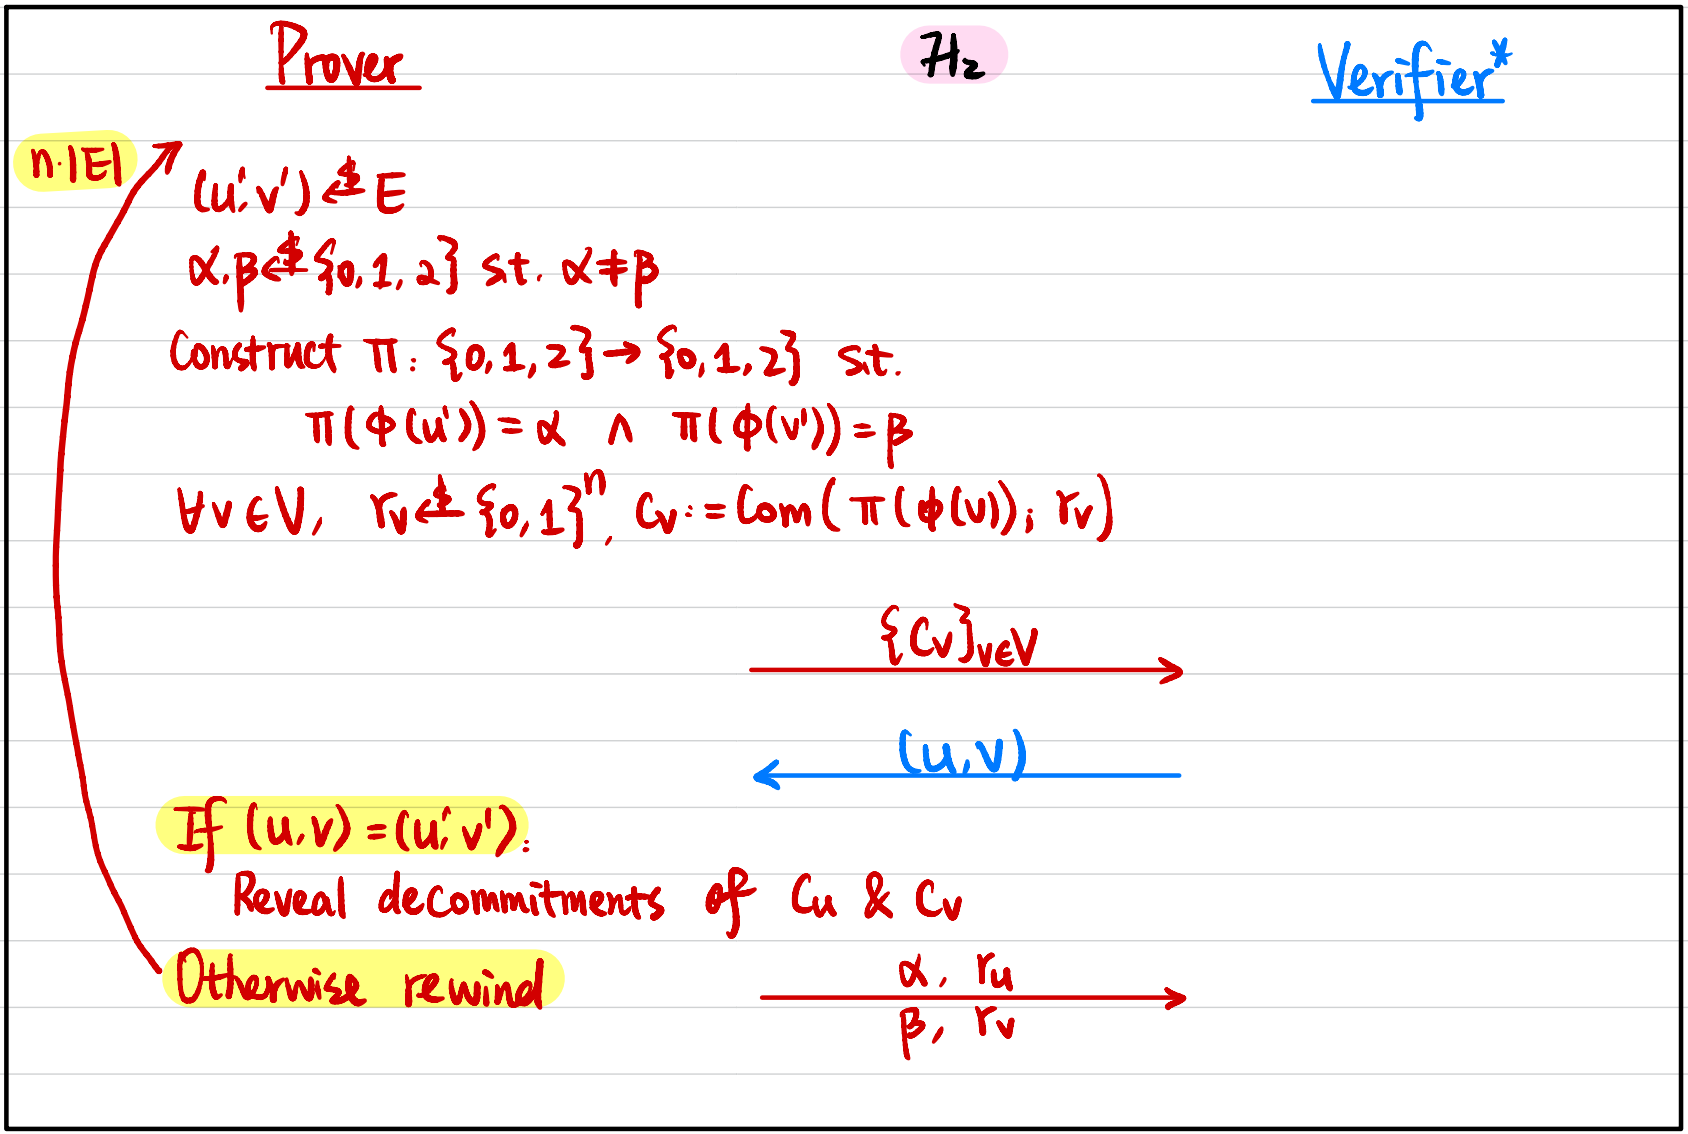

$$\Pr[\text{failure}] \leq \left(1 - \frac{1}{|E|}\right)^{n \cdot |E|} \simeq e^{-n}.$$

$\mathcal{H}_0$: Prover $\longleftrightarrow$ Verifier

$\equiv$ (identical distribution of π)

$\mathcal{H}_1$

$\approx_s$ (negligible failure probability)

$\mathcal{H}_2$

$\approx_c$ (computational hiding of Com)

$\mathcal{H}_3$: Simulator $\longleftrightarrow$ Verifier

Non-Interactive Zero-Knowledge (NIZK) Proof

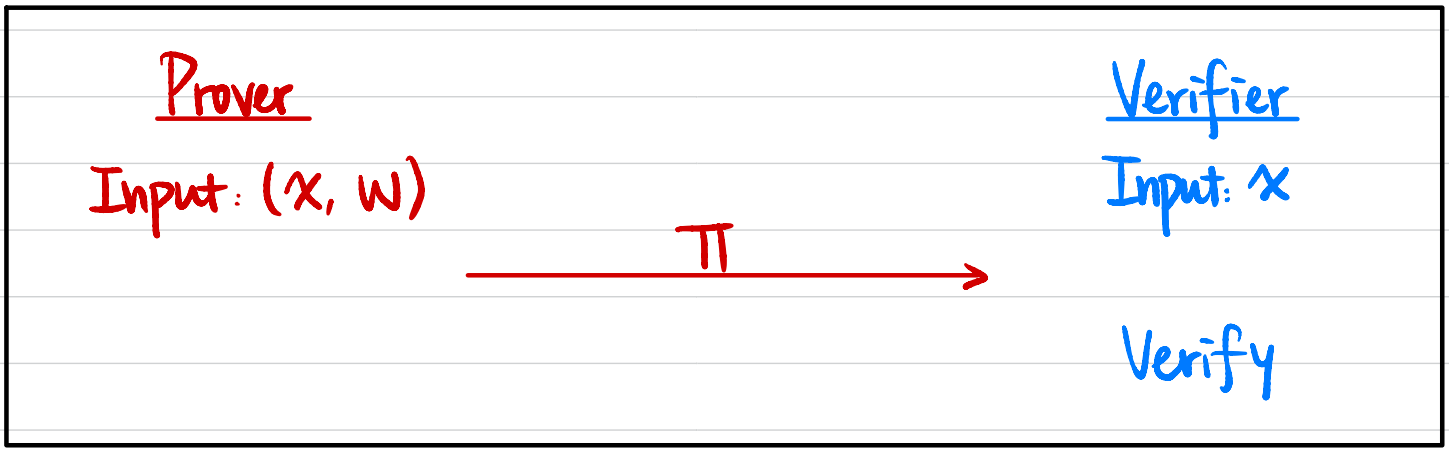

- **Completeness:** $\forall (x,w) \in R_L, \quad Pr[\, P(x,w) \longrightarrow V(x) \text{ outputs } 1\,] = 1.$
- **Soundness:** $\forall x \notin L, \ \forall P^*, \quad Pr[\, P^*(x) \longrightarrow V(x) \text{ outputs } 1\,] \leq negl(n)$
- **Zero-Knowledge:** $\forall PPT\ V^*, \ \exists PPT\ S \quad s.t. \quad \forall (x,w) \in R_L,$

$$Output_{V^*}[P(x,w) \longrightarrow V^*(x)] \simeq S(x)$$

Is it possible?

NOT in the "plain" model (assuming $P \neq NP$)

$x \in L$? $S(x) \rightarrow$ Verify

Model 1: Common Random String / Common Reference String (CRS)

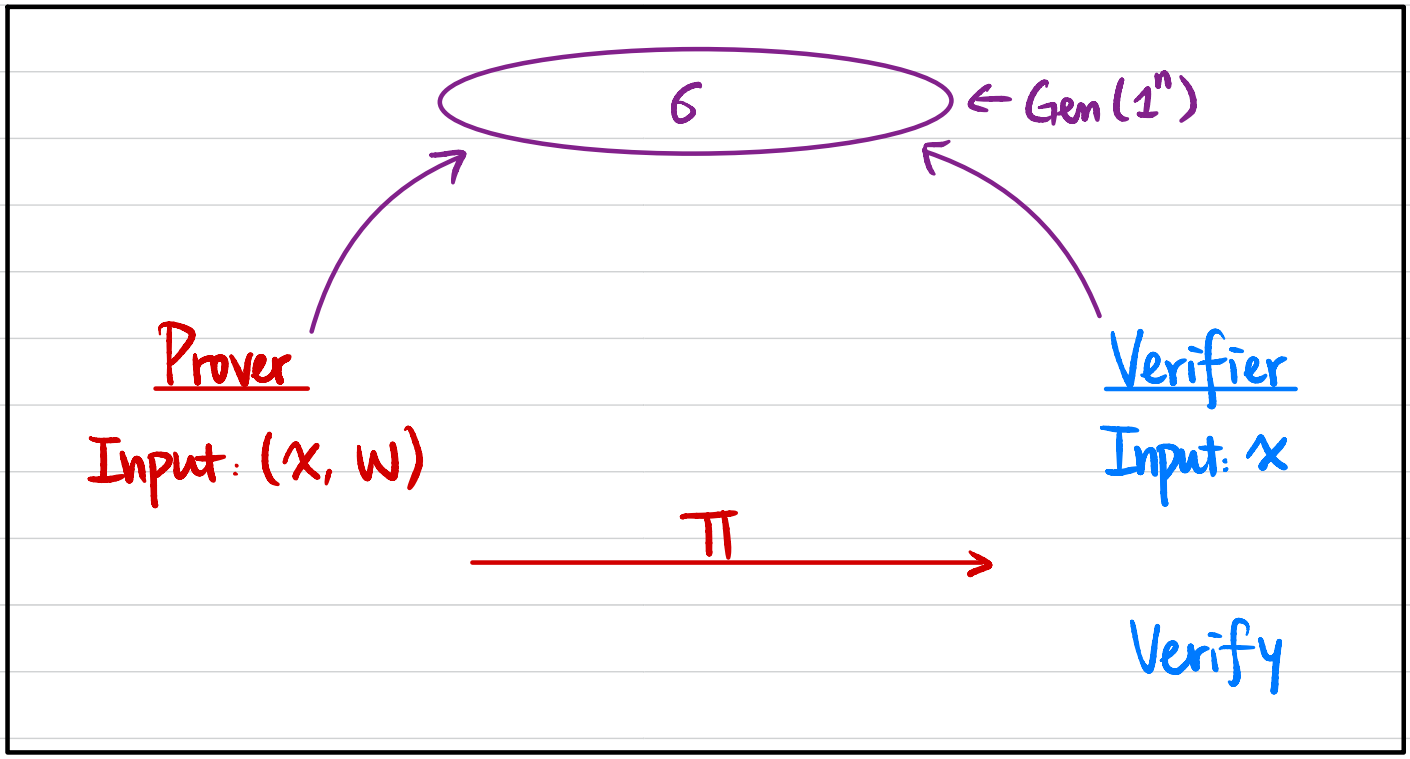

$S(x)$ generates both (σ, π)

- **Zero-Knowledge:** $\forall PPT\ V^*,\ \exists PPT\ S$ s.t. $\forall (x, w) \in R_L$,

$$Output_{V^*}\left[\sigma \leftarrow Gen(1^n),\ P(x, w, \sigma) \longrightarrow V^*(x, \sigma)\right] \simeq S(x)$$

Alternatively: $\left(\sigma \leftarrow Gen(1^n),\ P(x, w, \sigma)\right) \simeq S(x)$

Model 2: Random Oracle Model

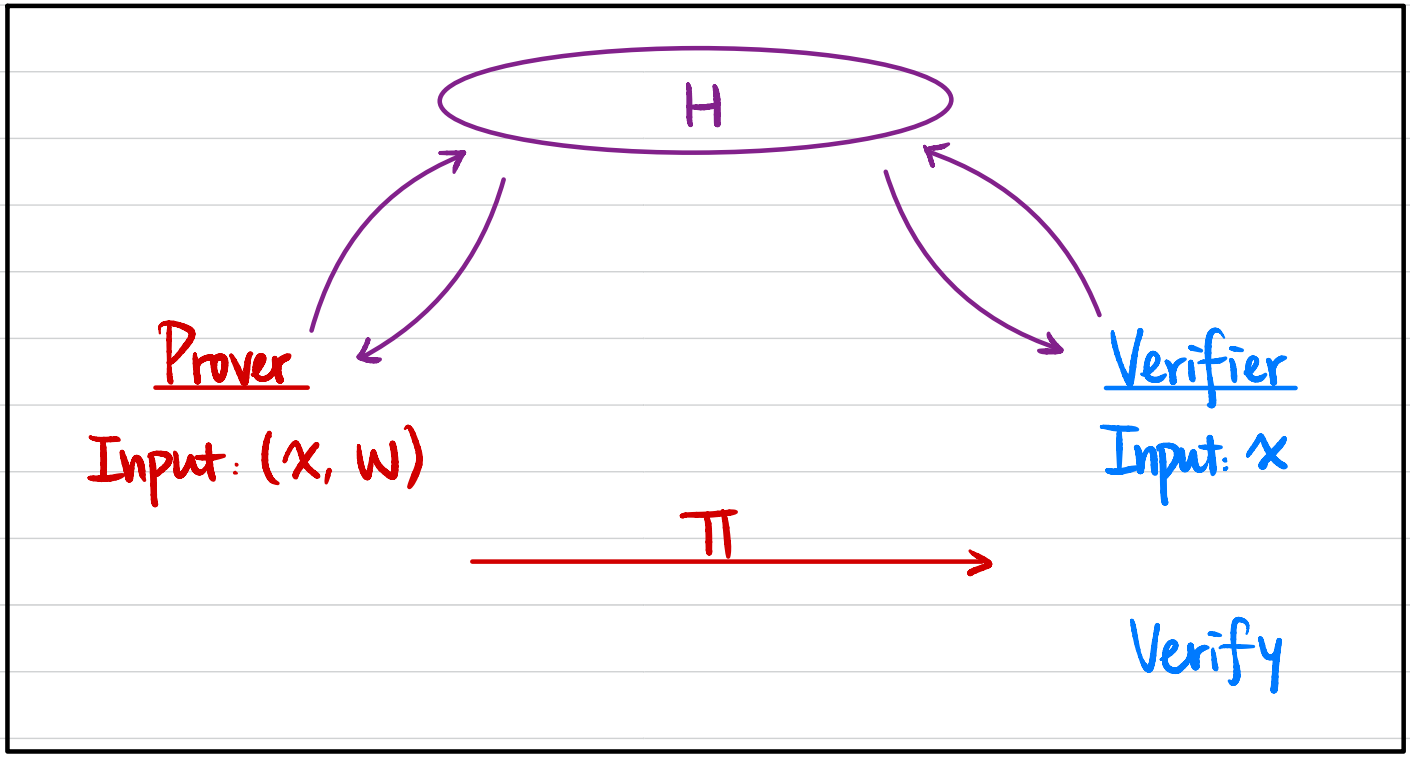

S controls input/output behavior of RO

Fiat-Shamir Heuristic

Public-Coin Honest-Verifier ZK (HVZK) $\Rightarrow$ NIZK in the RO model

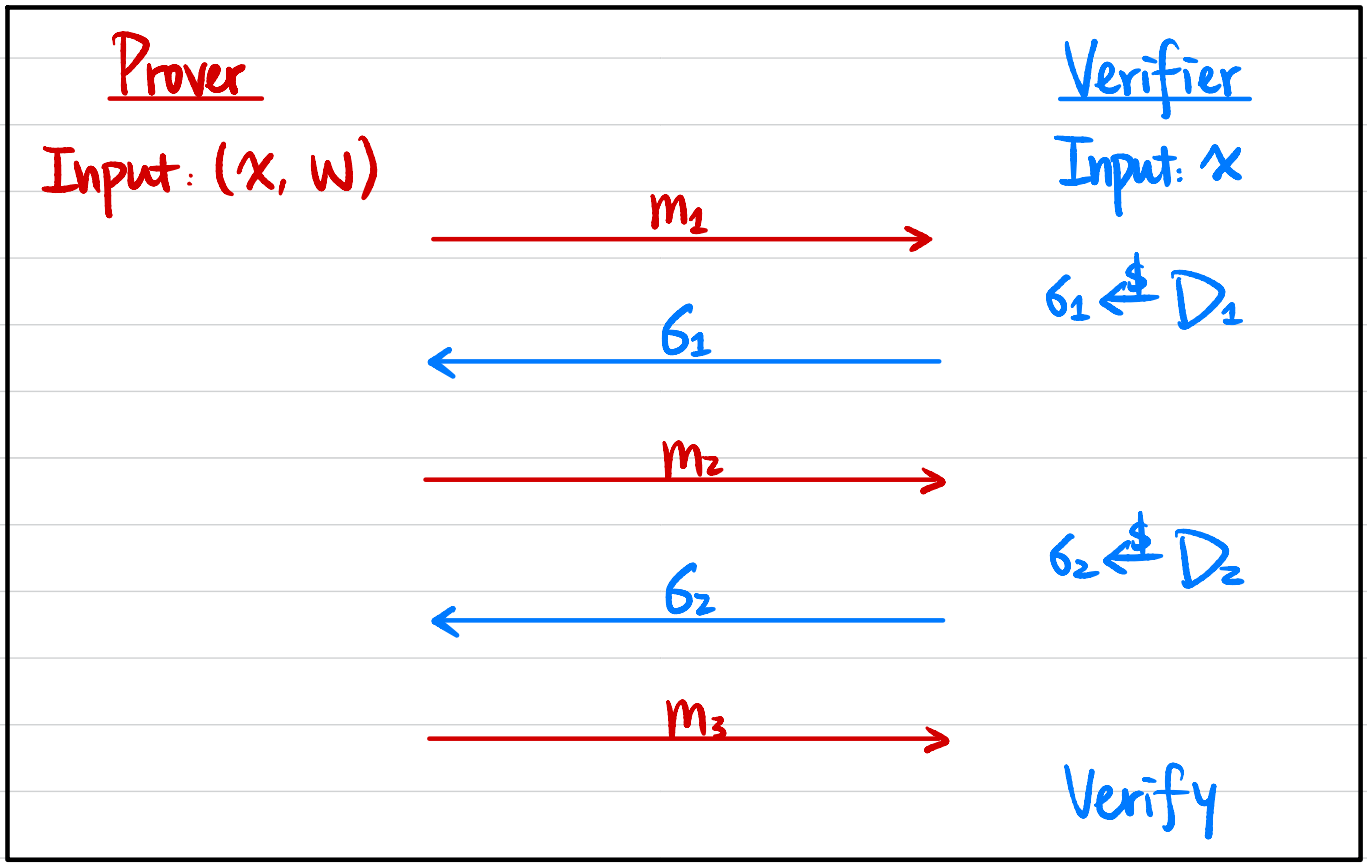

$\sigma_1 := H(x \| m_1)$

$\sigma_2 := H(x \| m_1 \| m_2)$

Secure Multi-Party Computation

Alice

x

Second date?

$f(x, y) = x \wedge y$

Bob

y

Who is richer?

$$f(x,y) = \begin{cases} 0 & \text{if } x > y \\ 1 & \text{otherwise} \end{cases}$$

Common friends?

$f(X, Y) = X \cap Y$

Secure Two-Party Computation (2PC)

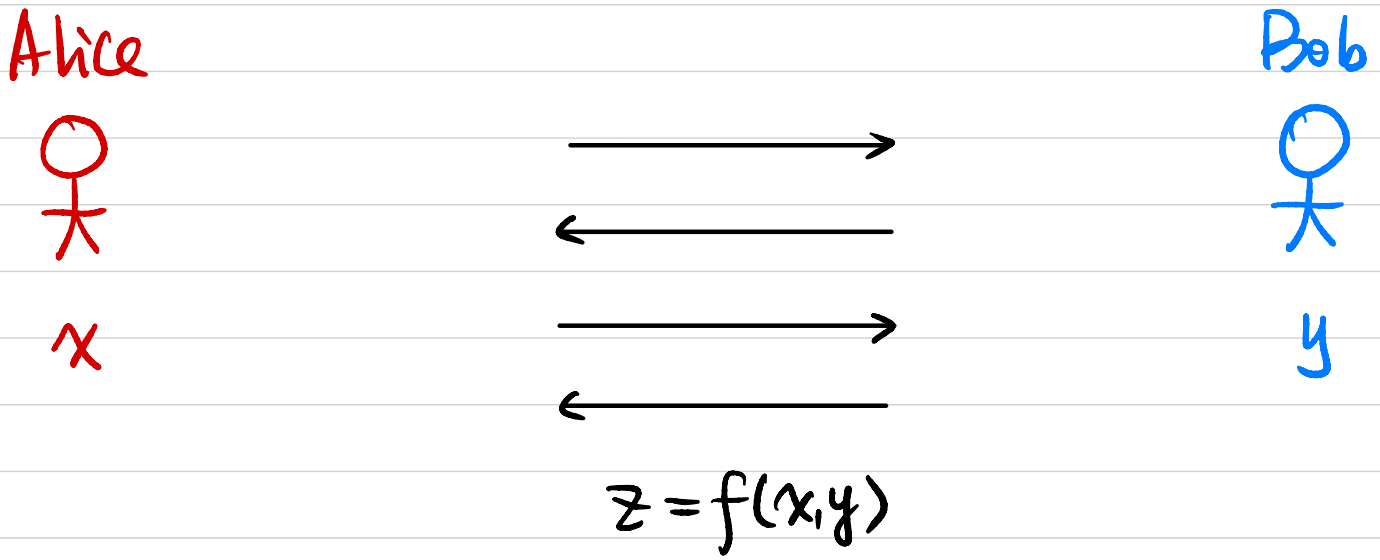

Applications:

- Password Breach Alert (Chrome / Firefox / Azure / iOS Keychain)
- Privacy-Preserving Contact Tracing for COVID-19 (Apple & Google)
- Ads Conversion Measurements / Personalized Advertising (Google / Meta)

Secure Multi-Party Computation (MPC)

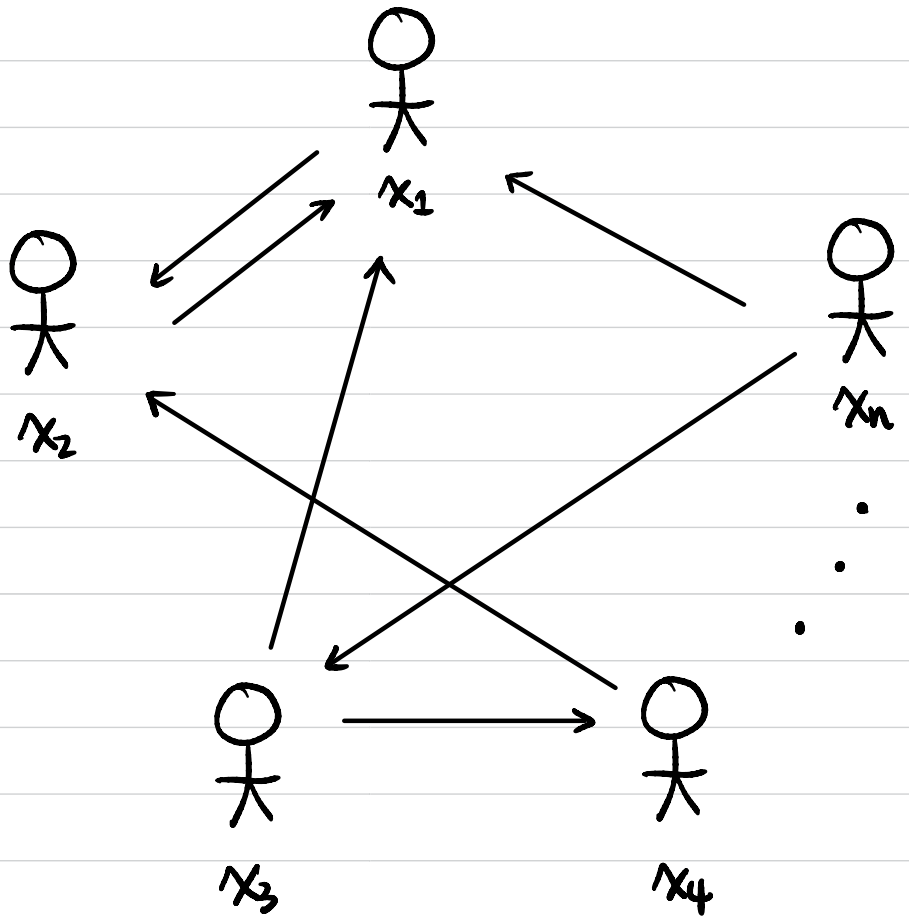

$$z = f(x_1, \cdots, x_n)$$

Secure Multi-Party Computation (MPC)

Applications:

- Privacy-Preserving Inventory Matching (J.P. Morgan)
- Setup Ceremony to securely generate CRS (Zcash)
- Distributed Key Management (Unbound / Coinbase)
- Federated Learning (Google Keyboard Search Suggestion)
- Auctions (Danish sugar beet auction)
- Boston gender wage gap (Boston Women's Workforce Council)

- Study / Analysis on Medical Data
- Fraud Detection (banks)

Setting

- n parties $P_1, P_2, \cdots, P_n$
  with private inputs $x_1, x_2, \cdots, x_n$

- Jointly compute $f(x_1, x_2, \cdots, x_n)$

- Communication:
  Authenticated secure point-to-point channels between each pair (P_i, P_j)
  (sometimes also assume broadcast channel)

- The adversary can "corrupt" a subset of the parties
  (e.g. at most t parties)

What properties do we want?

General Security Properties

- **Correctness:** The function is computed correctly.
- **Privacy:** Only the output is revealed.
- **Independence of Inputs:** Parties cannot choose inputs depending on others' inputs.

- **Security with Abort:** Adversary may "abort" the protocol.
  (preventing honest parties from receiving the output)
- **Fairness:** If one party receives output, then all receive output.
- **Guaranteed Output Delivery (GOD):** Honest parties always receive output.

Adversary's Power

Allowed adversarial behavior:

- **Semi-honest** / passive / honest-but-curious:

  Follow the protocol description honestly,

  but try to extract more information by inspecting transcript.

- **Malicious** / active:

  Can deviate arbitrarily from the protocol description.

Adversary's Computing Power:

- **Unbounded computing power** ⇒ Information-Theoretic (IT) Security
- **PPT bounded** ⇒ Computational Security

Security Against Semi-Honest Adversaries

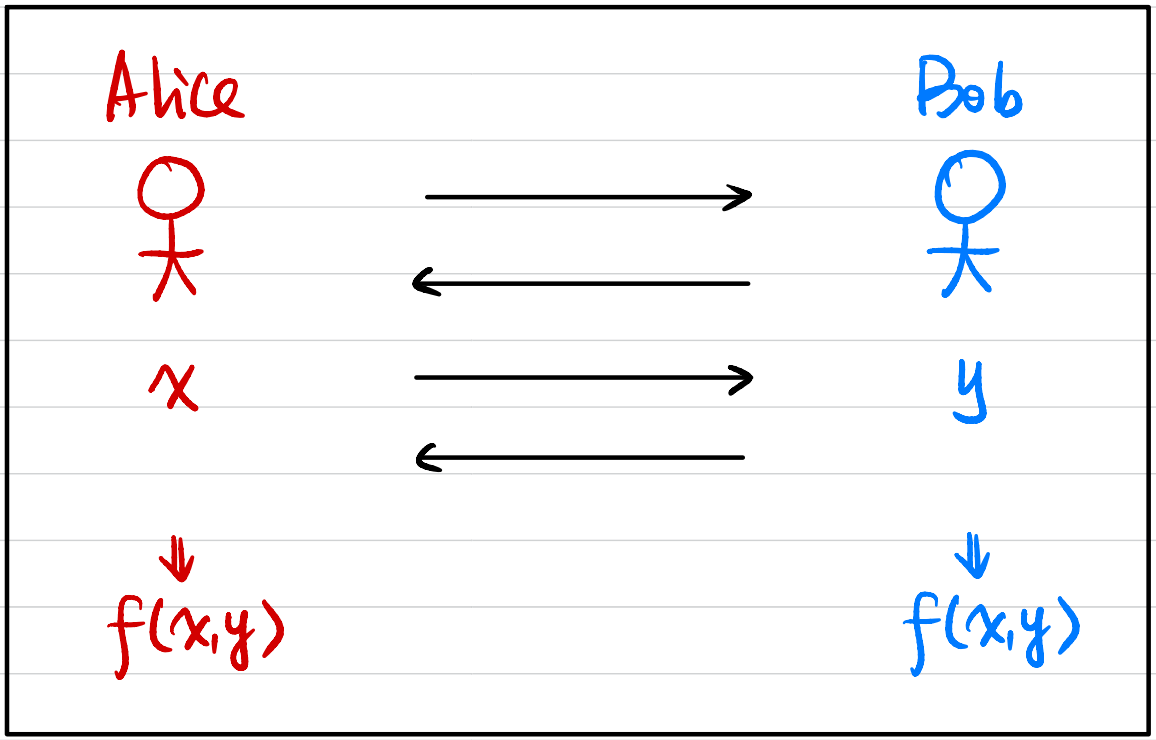

Alice's view:

$\text{View}_A^{\Pi}(x, y, n) := (x,$ internal random tape $r,$ messages from Bob$)$

Given x, $f(x,y)$, Alice's view can be "simulated".

Security Against Semi-Honest Adversaries

Def (Semi-honest security for 2PC)

Let f be a functionality. We say a protocol Π securely computes f against semi-honest adversaries if $\exists$ PPT algorithms S_A, S_B s.t. $\forall x, y$,

$$\left\{\begin{pmatrix} S_A(1^n, x, f(x,y)) \\ f(x,y) \end{pmatrix}\right\}_{n\in\mathbb{N}} \simeq \left\{\begin{pmatrix} \text{View}_A^{\Pi}(x,y,n) \\ \text{Output}^{\Pi}(x,y,n) \end{pmatrix}\right\}_{n\in\mathbb{N}}$$

$$\left\{\begin{pmatrix} S_B(1^n, y, f(x,y)) \\ f(x,y) \end{pmatrix}\right\}_{n\in\mathbb{N}} \simeq \left\{\begin{pmatrix} \text{View}_B^{\Pi}(x,y,n) \\ \text{Output}^{\Pi}(x,y,n) \end{pmatrix}\right\}_{n\in\mathbb{N}}$$

perfect / statistical / computational

$\equiv$ $\overset{s}{\simeq}$ $\overset{c}{\simeq}$

Security Against Malicious Adversaries

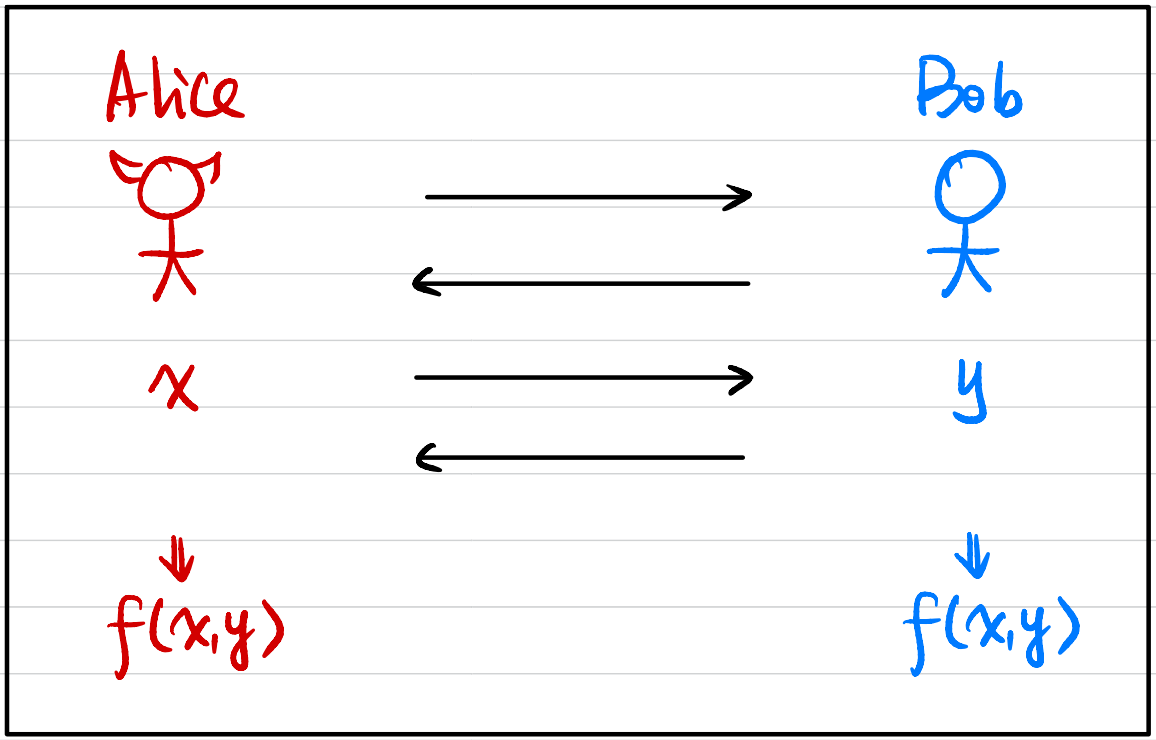

Alice's view:

$\text{View}_A^{\Pi}(x, y, n) := (x,$ internal random tape $r,$ messages from Bob$)$

Given x, $f(x,y)$, Alice's view can be "simulated".

↑ What output?

What's the best we can hope for? (Ideal World)

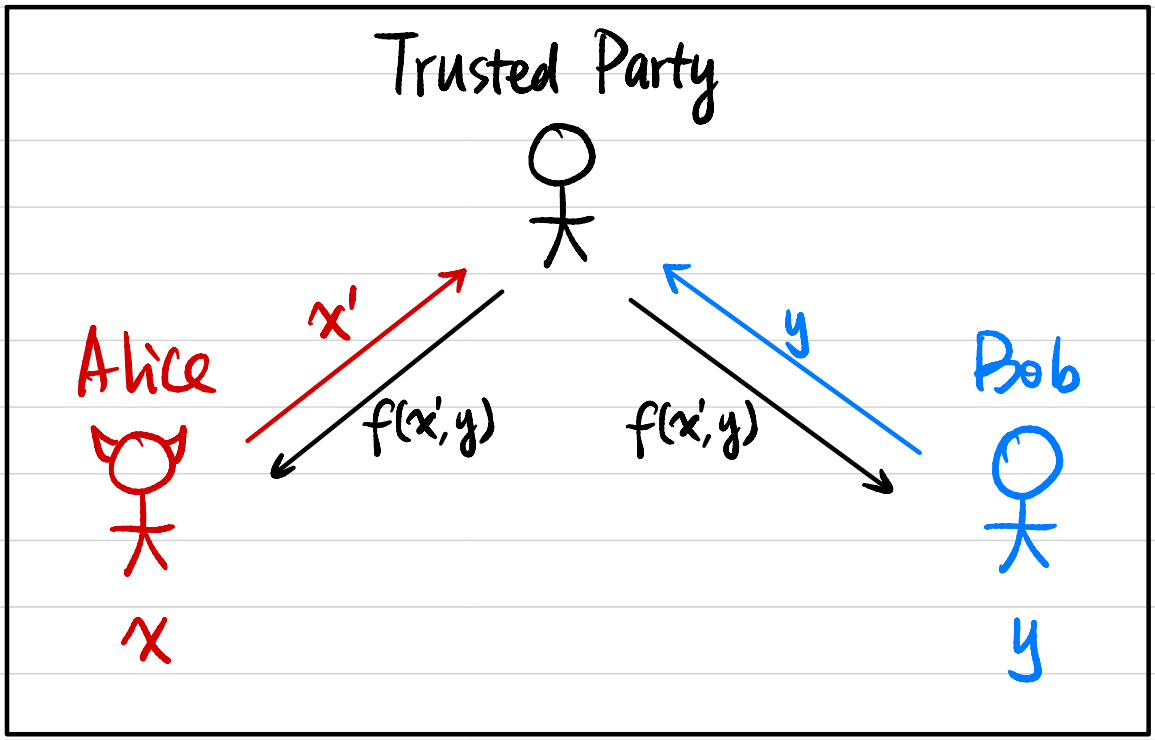

Security Against Malicious Adversaries (Real / Ideal Paradigm)

Execution in the Real World:

(PPT) adversary $\mathcal{A}$ corrupting party $i \in \{\text{Alice}, \text{Bob}\}$

$$\text{REAL}^{\pi}_{\mathcal{A},i} := \begin{pmatrix} \mathcal{A}\text{'s output} \\ \text{Honest party's output in Real World} \end{pmatrix}$$

Execution in the Ideal World:

PPT adversary $\mathcal{S}$ corrupting party $i \in \{\text{Alice}, \text{Bob}\}$

$$\text{IDEAL}^{f}_{\mathcal{S},i} := \begin{pmatrix} \mathcal{S}\text{'s output} \\ \text{Honest party's output in Ideal World} \end{pmatrix}$$

Def (malicious security for 2PC)

Let f be a functionality. We say a protocol π securely computes f against malicious adversaries if $\forall$ (PPT) $\mathcal{A}$ in the real word, $\exists$ PPT $\mathcal{S}$ in the ideal world s.t. $\forall i \in \{\text{Alice}, \text{Bob}\}$, $\forall x, y$,

$$\left\{\text{REAL}^{\pi}_{\mathcal{A},i}(x, y, n)\right\}_{n \in \mathbb{N}} \simeq \left\{\text{IDEAL}^{f}_{\mathcal{S},i}(x, y, n)\right\}_{n \in \mathbb{N}}$$