# Medical Data Privacy and Ethics in the Age of Artificial Intelligence

# Lecture 15: Differential Privacy and Access Control

Zhiyu Wan, PhD (wanzhy@shanghaitech.edu.cn)
Assistant Professor of Biomedical Engineering
ShanghaiTech University

April 30, 2025

# Homework 1 Hints

**Question 1a (8 points)**. According to [1], how many 3-digit ZCTAs are completely identifiable for this population (i.e., the entire population is expected to be uniquely identified) according to day of birth? Provide documentation on how you arrived at this result.

**Question 1b (8 points)**. According to [2], what are the 10 ZCTAs with the greatest proportion of uniquely identified individuals based on their day of birth? Provide documentation on how you arrived at this result.

**Question 1c (9 points)**. Provide a scatterplot of the unique identifiability estimates for 3-digit ZCTAs according to day of birth for [1] vs [2]. At what point do the estimates from [1] significantly differ from [2]?

- **Question 1d (9 points)**. Repeat the analysis in Question 1c, but now compute the expected proportion of the populations that are in a group of size 2 or smaller. At what point do the estimates from [1] significantly differ from [2]?

# Randomized (Golle '06)

- Don't always have exact knowledge of what a data recipient has access to
  - Disclose sample with {*dob, gender, zip*}, but don't know the population's values
- May know population counts, such as
  - U.S. Census aggregates for {*year of birth, gender, county*}
- Conversion: {Year of Birth, ZIP} → {Date of Birth, ZIP}
- Alternative option: Randomly allocate 12,000 "people" to 365 cells

| | | Birth Year |
|---|---|---|
| | | 1980 |
| ZIP | zip1 | 12000 |
| | zip2 | 50000 |
| | … | |
| | zip m | 10000 |

| | | Birthdate | | | SUM |
|---|---|---|---|---|---|
| | | 1/1/80 | … | 12/31/80 | |
| ZIP | zip1 | random | | random | 12000 |
| | zip2 | random | | random | 50000 |
| | … | | | | |
| | zip m | random | | random | 10000 |

# It's an Occupancy Problem (Golle '06)

- $n$ people in aggregated bin
- $b$ disaggregated bins
- the expected # of bins with exactly $i$ people

$$f_i(n) = \binom{n}{i} b^{1-n} (b-1)^{n-i}$$

- Total number of people in a group of size less than $k$

$$r_k(n) = \sum_{i=1}^{k-1} f_i(n)$$

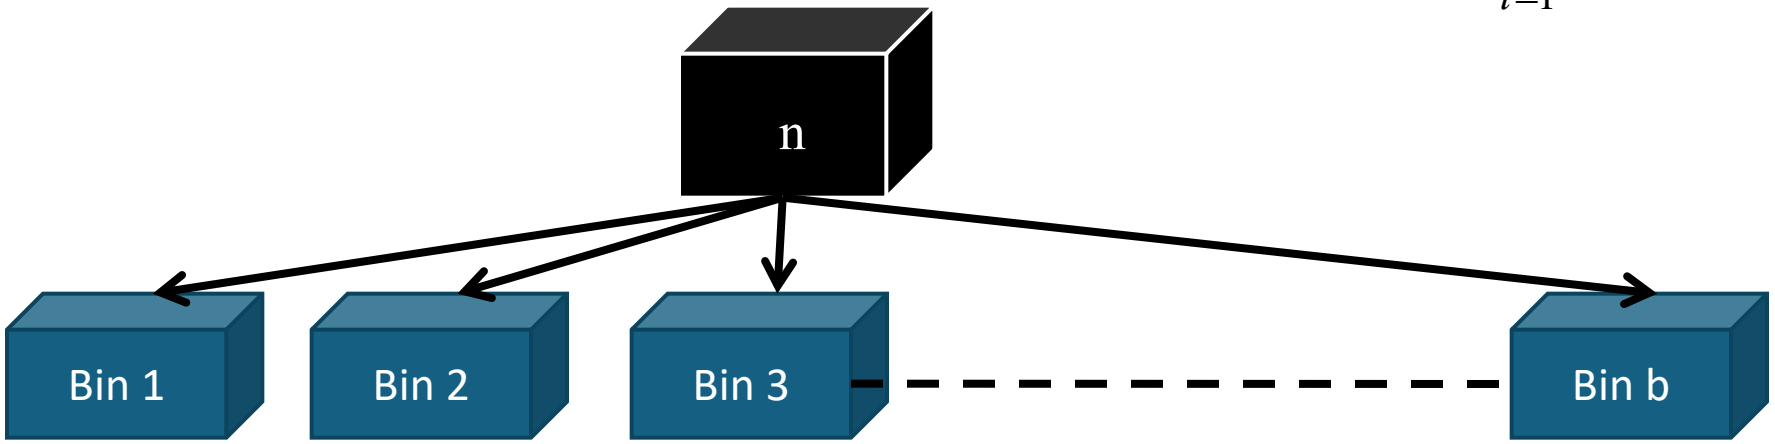

# Sample Calculation

| | | *Quasi-ID values (Bins)* | | | | | |
|---|---|---|---|---|---|---|---|
| | | **2** | **4** | **256** | **512** | **1024** | **8192** |
| **Population (Balls)** | **2** | 0.5 | 2.25 | 254.01 | 510.00 | 1022.00 | 8190.00 |
| | **4** | 0.125 | 1.26 | 252.03 | 508.02 | 1020.01 | 8188.00 |
| | **64** | $1.08 \times 10^{-19}$ | $4.04 \times 10^{-8}$ | 199.37 | 451.84 | 961.96 | 8128.25 |
| | **1024** | 0.00 | 0.00 | 4.69 | 69.29 | 376.71 | 7229.41 |
| | **2048** | 0.00 | 0.00 | 0.09 | 9.38 | 138.58 | 6379.94 |

Expected Number of Quasi-ID values with 0 people

# Sample Calculation

Quasi-ID values (Bins)

| Population (Balls) | 2 | 4 | 256 | 512 | 1024 | 8192 |
|---|---|---|---|---|---|---|
| **2** | 0.250 | 0.563 | 0.992 | 0.996 | 0.998 | 1.000 |
| **4** | 0.063 | 0.315 | 0.984 | 0.992 | 0.996 | 1.000 |
| **64** | 0.000 | 0.000 | 0.779 | 0.883 | 0.939 | 0.992 |
| **1024** | 0.000 | 0.000 | 0.018 | 0.135 | 0.368 | 0.882 |
| **2048** | 0.000 | 0.000 | 0.000 | 0.018 | 0.135 | 0.779 |

Expected Number of  Quasi-ID values with 0 people

# Birthday Problem

- Assume birthday is uniformly distributed at random over the year.
- If $n$ people are born in a year, **the expected number of days on which exactly 1 person born** is

$$f_1(n) = n * \left(\frac{364}{365}\right)^{n-1}$$

# Golle's Approach

- Special case of general equation
- If *n* people are born in a year, **the expected # of days on which exactly *k* people born** is

$$f_k(n) = \binom{n}{k} 365^{1-n} 364^{n-k}$$

# Hints

- **Q1a**
- According the paper of Sweeney's, we should assume the day of birth is uniformly distributed within each age group (actually assuming the number of birth for every day is exactly equals to the expected number of birth every day!). For simplicity, I use age=100 as a maximum.
  First compute the thresholds (possible values) for each age group

| Group | Start Age | End Age | #years | * | days per year | = | thresholds |
|---|---|---|---|---|---|---|---|
| 1 | 0 | 4 | 5 | | 365 | | 1825 |
| 2 | 5 | 9 | 5 | | 365 | | 1825 |
| 3 | 10 | 14 | 5 | | 365 | | 1825 |
| 4 | 15 | 17 | 3 | | 365 | | 1095 |
| 5 | 18 | 19 | 2 | | 365 | | 730 |
| 6 | 20 | 20 | 1 | | 365 | | 365 |
| 7 | 21 | 21 | 1 | | 365 | | 365 |
| 8 | 22 | 24 | 3 | | 365 | | 1095 |
| 9 | 25 | 29 | 5 | | 365 | | 1825 |
| 10 | 30 | 34 | 5 | | 365 | | 1825 |

For every 3-Digit ZCTA, use these twenty-three thresholds to minus twenty-three corresponding population, if all twenty-three age groups get positive number this ZCTA is marked as completely identifiable.

# Hints

- **Q1b**
- The number of individuals that can be k-uniquely identified in certain region and certain time periods is as follows:

$$f_k(n) = \binom{n}{k} N^{1-n} (N-1)^{n-k}$$

$$= \binom{n}{k} N^{1-k} \left(\frac{N-1}{N}\right)^{n-k}$$

Where n is population in this region, N is number of days in this time period.

The computation for identified populations in group1 in 006 ZCTA is as follows:

k=1; n=42864; N=1825;

$$f_k(n) = \binom{n}{k} N^{1-k} \left(\frac{N-1}{N}\right)^{n-k}$$

$$= \binom{42864}{1} 1825^{0} \left(\frac{1824}{1825}\right)^{42864}$$

$$= 2.6866250497 \times 10^{-6}$$

# Homework 1 Hints

- Question 2

- **Question 2a (17 points).** Use the one-sided t-test method proposed by Homer et al. [3] with a 90% confidence level. According to this method, which of the target individuals are predicted to be in the reference population? Show your work. Note 1: there are 9 degrees of freedom, so the 90% confidence level in this computation corresponds to a t-value of **1.38** or greater. Note 2: Remember, 00 contributes +0; "01" contributes +0.5.; "11" contributes +1.

- **Question 2b (17 points).** Using the one-sided t-test method described by Homer et al [3], which of the individuals predicted to be in the study, are predicted to be in the positive and negative diseases classes respectively? **Show your work**.

***Programming Code***

- One-sided t-test

# Homer's Attack in a Nutshell

## The attacker knows:

- The genome of the target (her set of genomic variants) - $Y_{ij}$
- The allele frequencies of the Mixture he's attacking - $M_j$
- Population allele frequencies - $Pop_j$

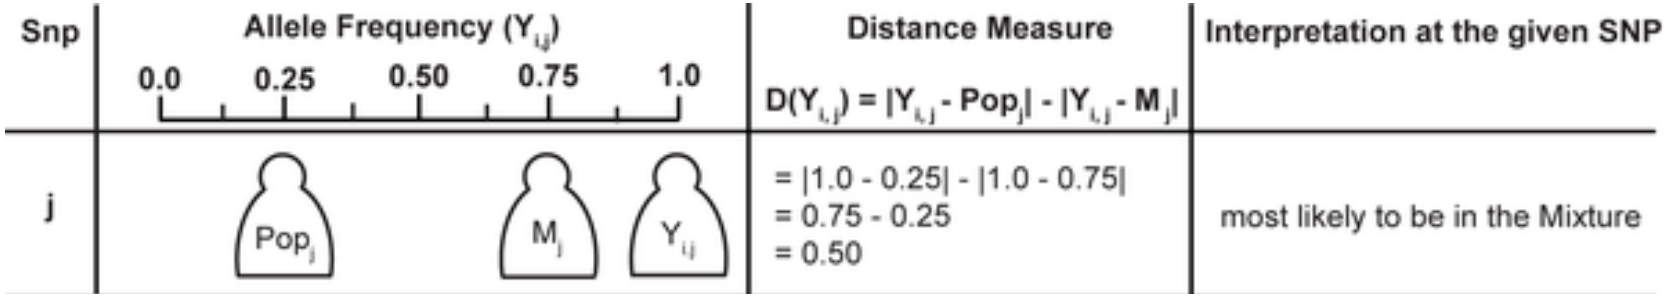


Figure from: Homer N, et al. PLoS Genetics. 2008; 4(8): e1000167.

# So, where's the Target?

- $|Y_{ij} - M_j|$ ← difference between individual & mixed study
- $|Y_{ij} - R_j|$ ← difference between individual & reference pop.

$$D(Y_{ij}) = |Y_{ij} - R_j| - |Y_{ij} - M_j|$$

- Null Hypothesis: Individual is not in mixed study.
  - $D(Y_{ij})$ should be approaching 0 [due to "ancestral similarity" in M and R]
- Alternative Hypothesis
  - $D(Y_{ij}) > 0$ because $M_j$ is shifted away from reference by $Y_j$'s contribution to the mixture
  - $D(Y_{ij}) < 0$ because $Y_j$ is more similar to reference population than the mixture

# Testing

$$T(Y_i) = \frac{E(D(Y_i)) - \mu_0}{SD(D(Y_i))/\sqrt{s}}$$

- $\mu_0$ : Mean of $D(Y_i)$ of all individuals **not** in the mixture
- $SD(Y_i)$: St. Dev. of $D(Y_{i,j})$ for all SNPs j and individual $Y_i$
- s: number of SNPs
- Can assume $\mu_0 = 0$ [random individual equidistant to M & R]
- Null hypotheses T = 0. Alternative is that T > 0 $(T > \theta = 1.38)$

# Hints

$$T(Y_i) = \frac{E(D(Y_i)) - \mu_0}{SD(D(Y_i))/\sqrt{s}}$$

- Q2a
  - Null Hypothesis: Individual is not in the reference.

$$D(Y_{ij}) = |Y_{ij} - M_j| - |Y_{ij} - R_j|$$

- Q2b
  - Null Hypothesis: Individual is not in the mixture.

$$D(Y_{ij}) = |Y_{ij} - R_j| - |Y_{ij} - M_j|$$

  - Null Hypothesis: Individual is not in the positive class

$$D(Y_{ij}) = |Y_{ij} - C_j^-| - |Y_{ij} - C_j^+|$$

  - Null Hypothesis: Individual is not in the negative class

$$D(Y_{ij}) = |Y_{ij} - C_j^+| - |Y_{ij} - C_j^-|$$

# Adjustment to HW1

**Question 3a (25 points)**: Generate a Manhattan plot and identify the SNP(s) significantly associated with each phenotype. (5 points for each phenotype.)

- **Question 3b (7 points)**: Provide the R code (or code in other programming language) in a single file along with your submission and ensure it runs correctly.

**Question 3a (25 points)**: Generate a Manhattan plot and identify the SNP(s) significantly associated with each phenotype. (**25** points for one phenotype. **Additional 5 points for each additional phenotype.**)

- **Question 3b (7 points)**: Provide the R code (or code in other programming language) in a single file along with your submission and ensure it runs correctly.

# Learning Objectives of This Lecture

After this lecture, students should be able to:

- Know the concept of differential privacy (DP)
- Know the concept of role-based access control (RBAC)

# Sanitization of Databases

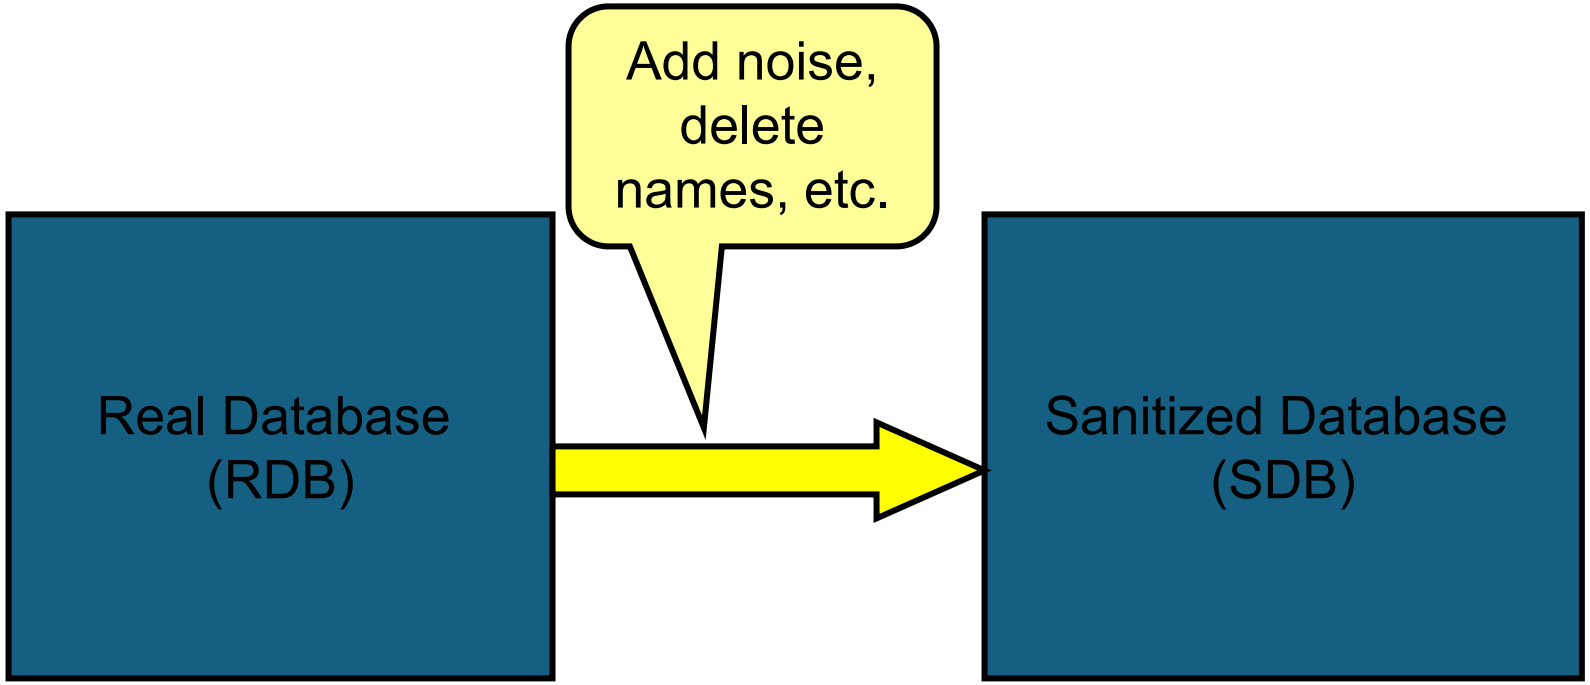


Health records

Census data

Protect privacy

Provide useful information (utility)

Adapted from Dr Malin's slides

# Basic Setting

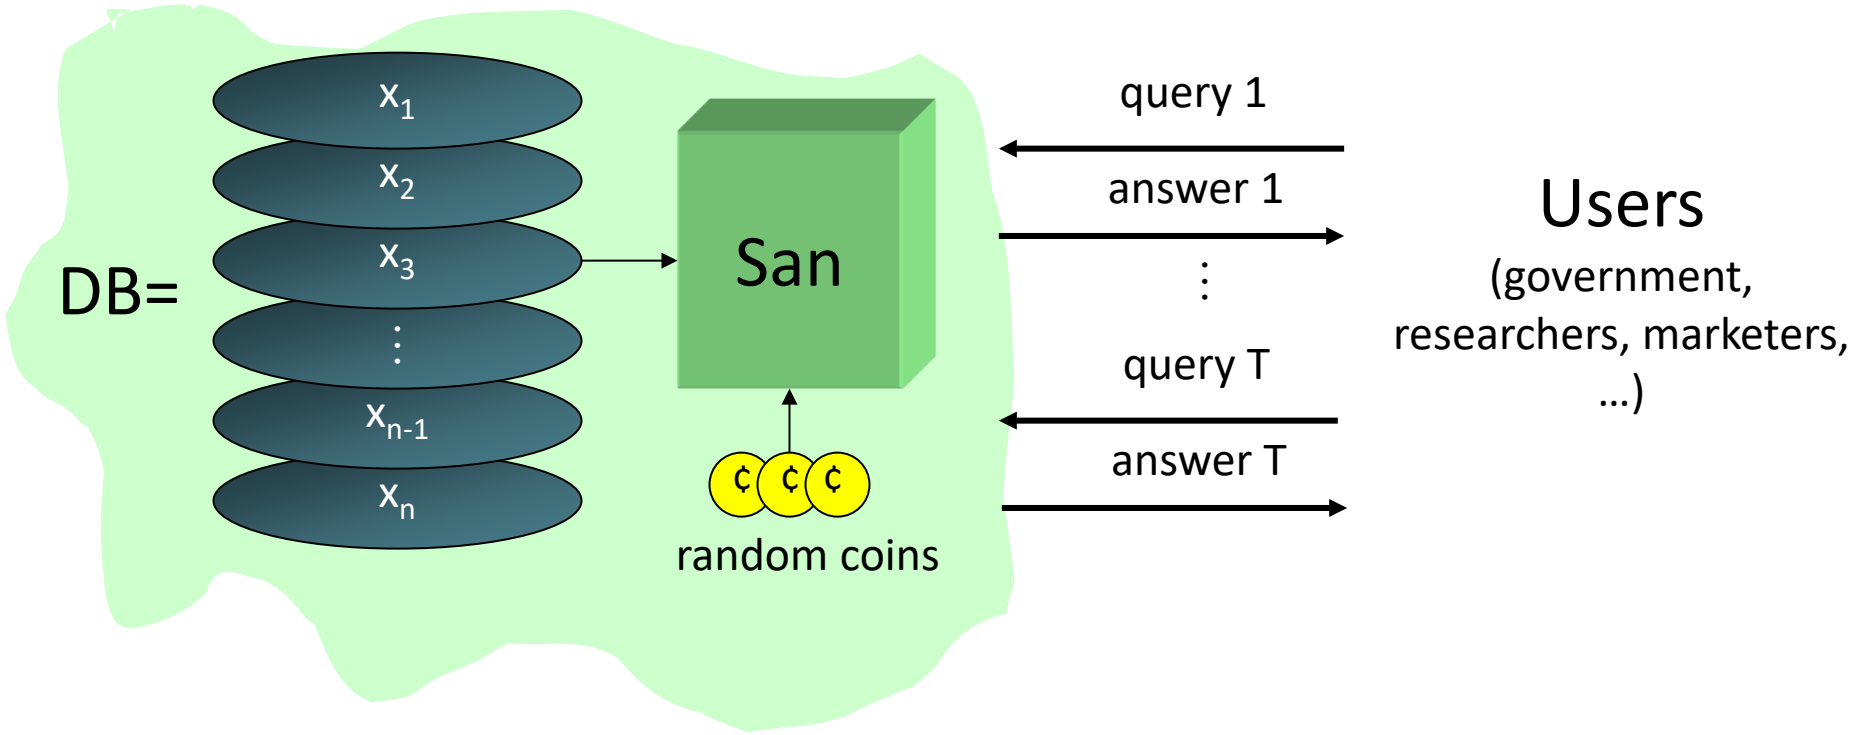

# Examples of Sanitization Methods

- Input perturbation
  - Add random noise to database, release
- Summary statistics
  - Means, variances
  - Marginal totals
  - Regression coefficients
- Output perturbation
  - Summary statistics with noise
- Interactive versions of the above methods
  - Auditor decides which queries are OK, type of noise

# Differential Privacy (informal)

Output is similar whether  any si
is included in the database or no

If there is already some risk of revealing a secret of C by combining auxiliary information and something learned from DB

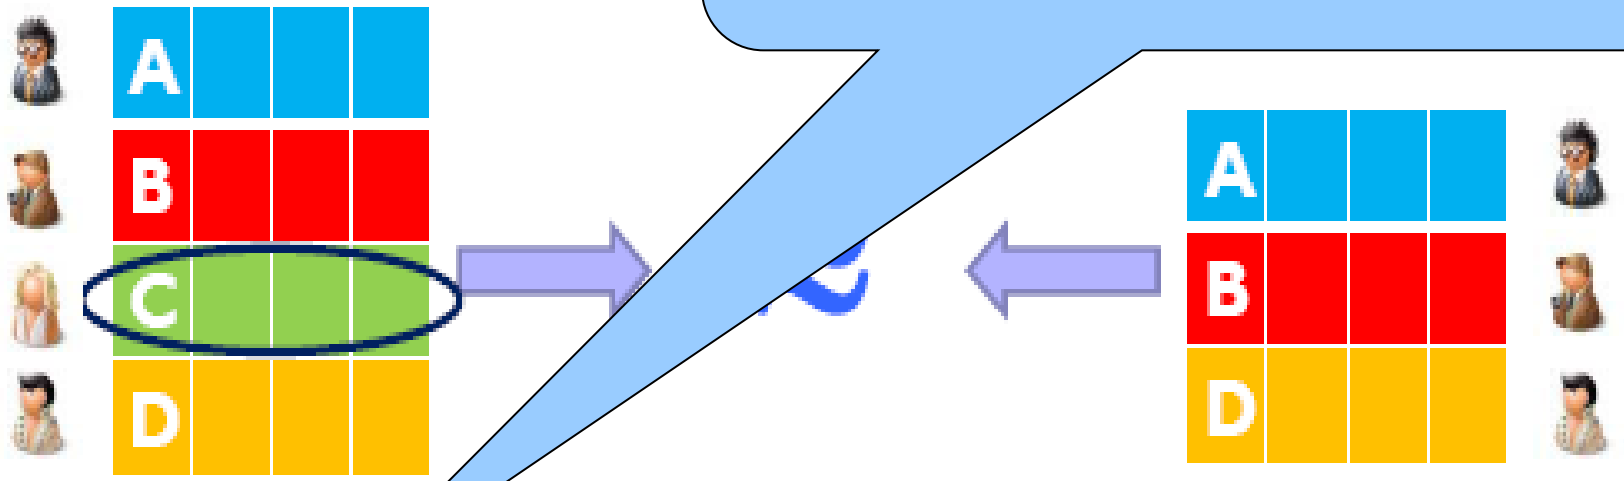


C is no worse off  because her record is included in the computation

# Differential Privacy (informal)

Output is similar whether  any si
is included in the database or no

If there is already some risk of revealing a secret of C by combining auxiliary information and something learned from DB, then that risk is still there

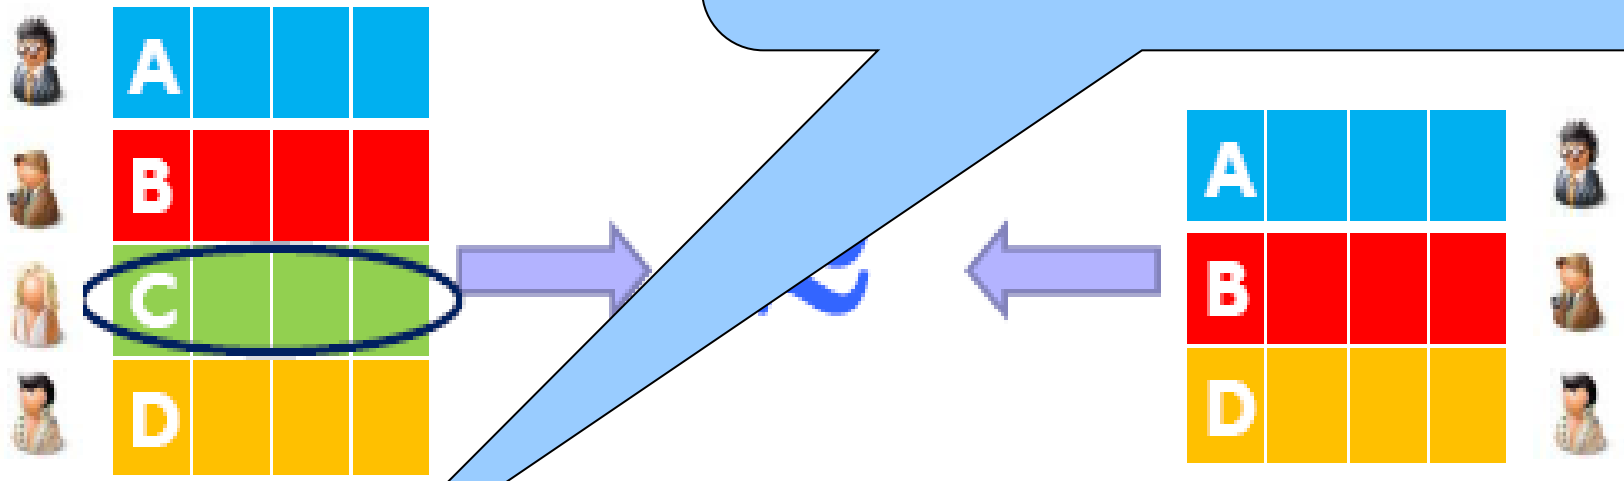


C is no worse off  because her record is included in the computation

# Differential Privacy (informal)

Output is similar whether  any si
is included in the database or no

If there is already some risk of revealing a secret of C by combining auxiliary information and something learned from DB, then that risk is still there but not increased by C's participation in the database

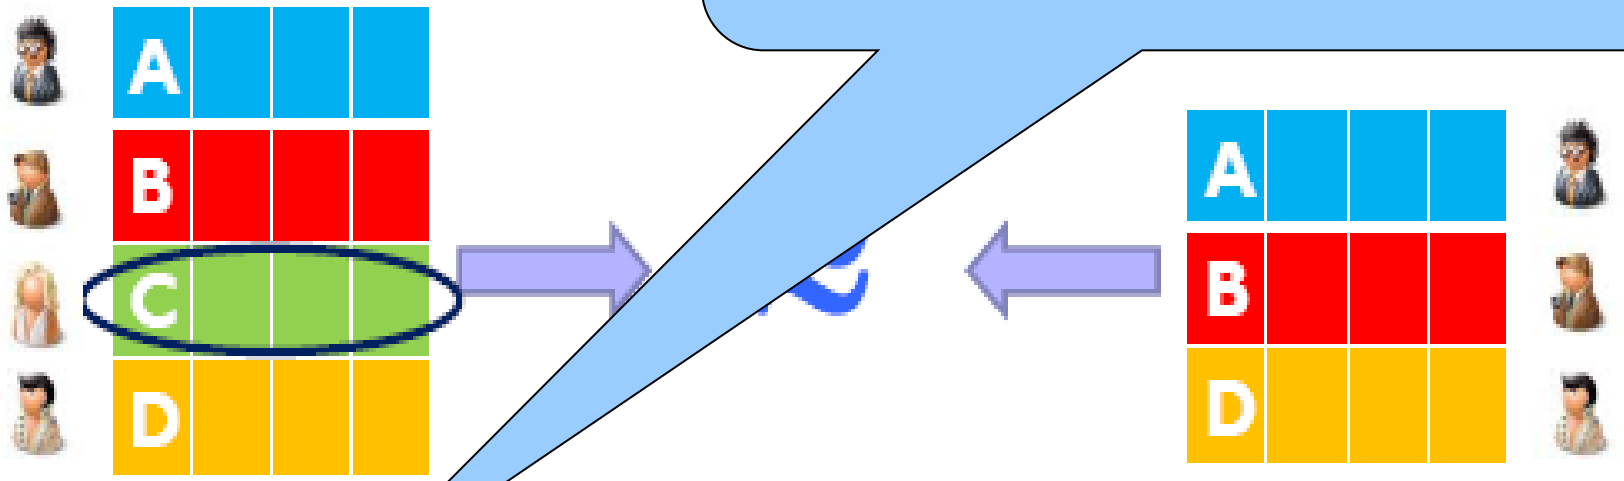


C is no worse off  because her record is included in the computation

# Differential Privacy is …

- … a guarantee about statistical confidentiality
  - The behavior of the system -- probability distribution on outputs -- is essentially unchanged, independent of whether any individual opts in or opts out of the dataset

- … a type of indistinguishability of behavior on neighboring inputs
  - Suggests other applications:
    - Approximate truthfulness as an economics solution concept (e.g., mechanism design)
    - As alternative to functional (or syntactic) privacy (e.g., *k-anonymity*)

- … useless without utility guarantees
  - Typically, “one size fits all” measure of utility
  - Simultaneously optimal for different priors, loss functions

# Strawman Definition

- Assume $x_1, \ldots, x_n$ are drawn i.i.d. (independent and identically distributed) from an unknown distribution
- Candidate definition: sanitization is safe if it only reveals the distribution
- Implied approach:
  - Learn the distribution
  - Release description of distribution or resample points

# Challenges with Classic Intuition

- Popular interpretation: prior and posterior views about an individual shouldn't change "too much"

- How much is "too much?"
  - Can't achieve small levels of disclosure and keep the data useful
  - Adversarial user is supposed to learn unpredictable things about the database

# Impossibility Result

[Dwork]

- Privacy: for some definition of "privacy breach,"

  $\forall$ distribution on databases, $\forall$ adversaries A, $\exists$ A′

  such that $Pr(A(San)=breach) - Pr(A'()=breach) \leq \varepsilon$

  - For reasonable "breach", if San(DB) contains information about DB, then some adversary breaks this definition

- Example
  - Alice knows that Bob is 2 inches taller than the average Male
  - DB allows computing average height of a Male
  - This DB breaks Bob's privacy according to this definition… even if his record is not in the database!

# Differential Privacy (1)

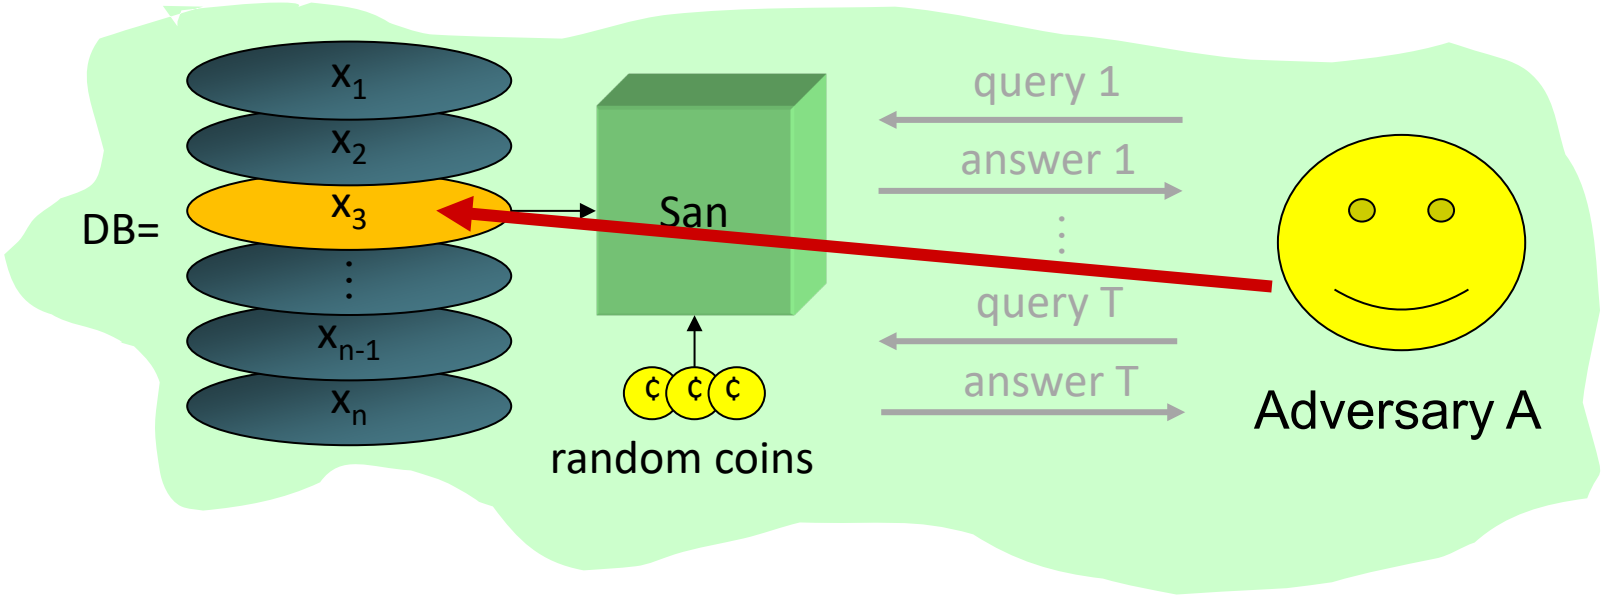


- Example with Males and Bob
  - Adversary learns Bob's height even if he is not in the database

- Intuition: "Whatever is learned would be learned regardless of whether or not Bill participates"
  - Dual: Whatever is already known, situation won't get worse

# Differential Privacy (2)

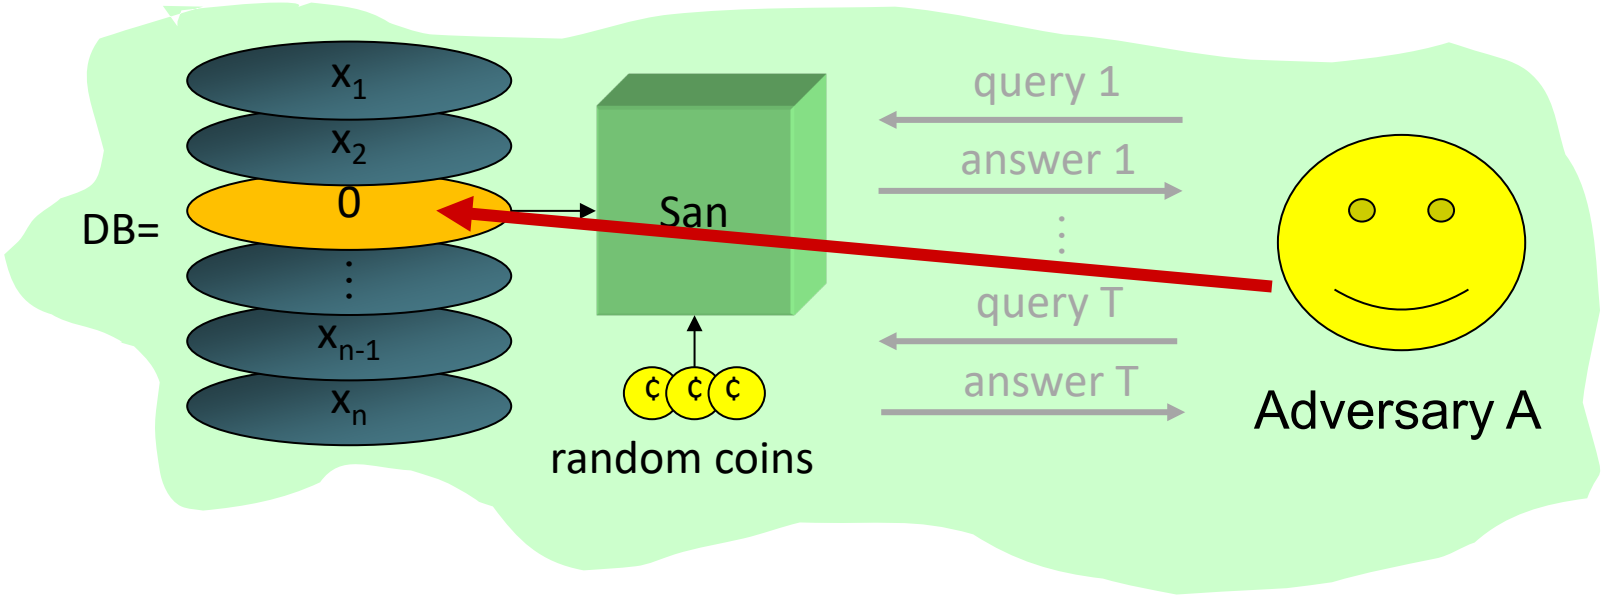


- Define n+1 games

# Differential Privacy (2)

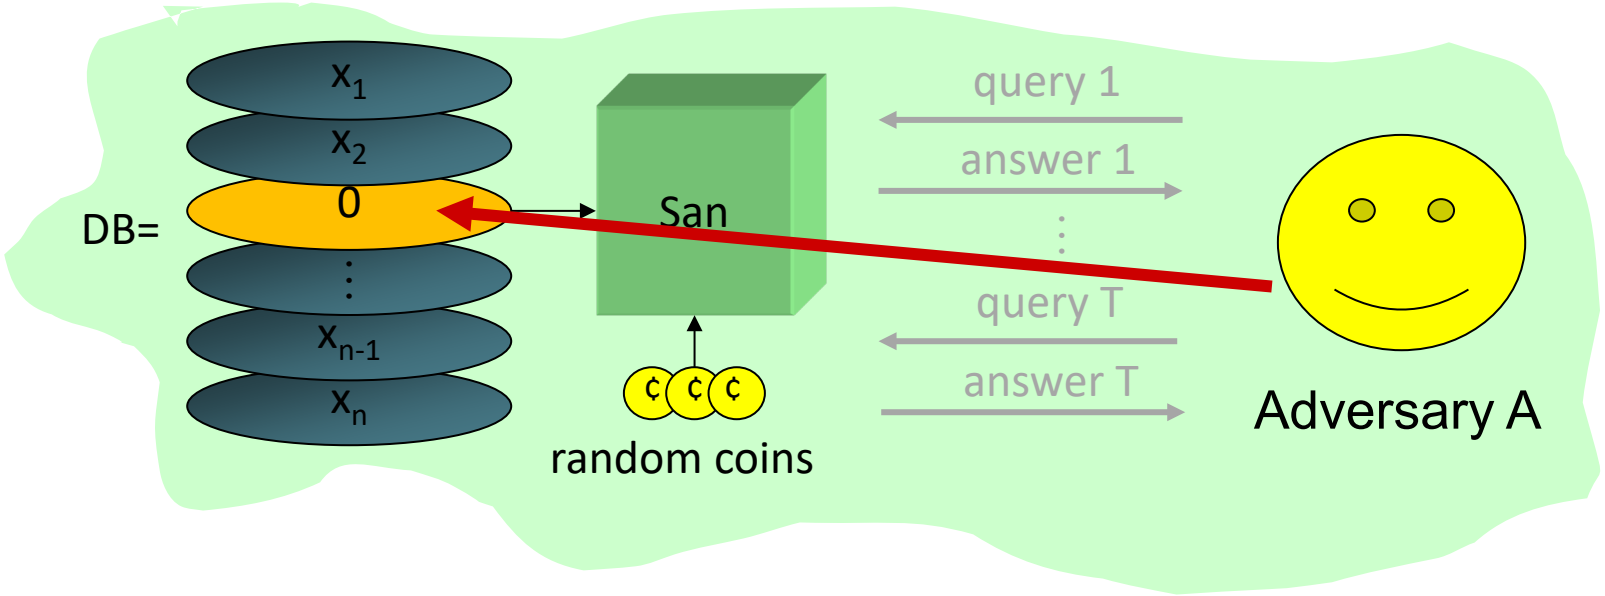


- Define n+1 games
  - Game 0: Adv. interacts with San(DB)

# Differential Privacy (2)

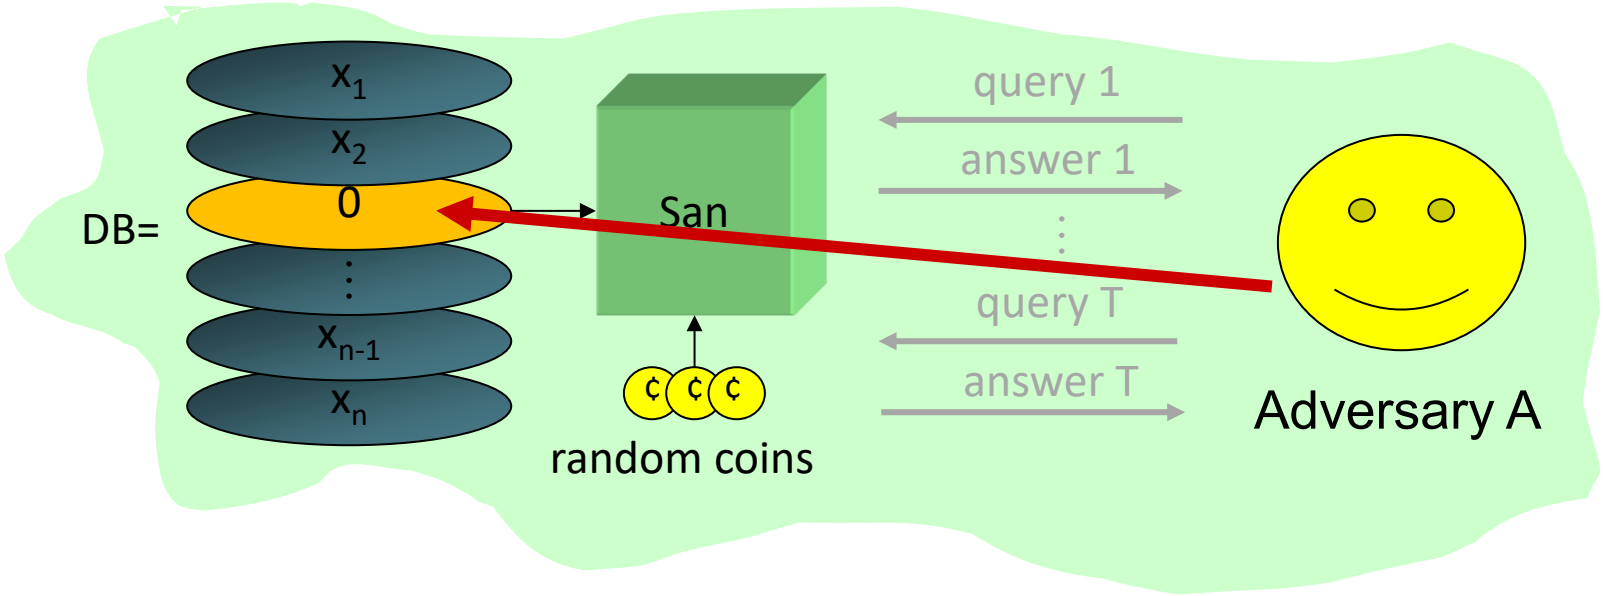


- Define n+1 games
  - Game 0: Adv. interacts with San(DB)
  - Game i: Adv. interacts with San($DB_{-i}$);

# Differential Privacy (2)

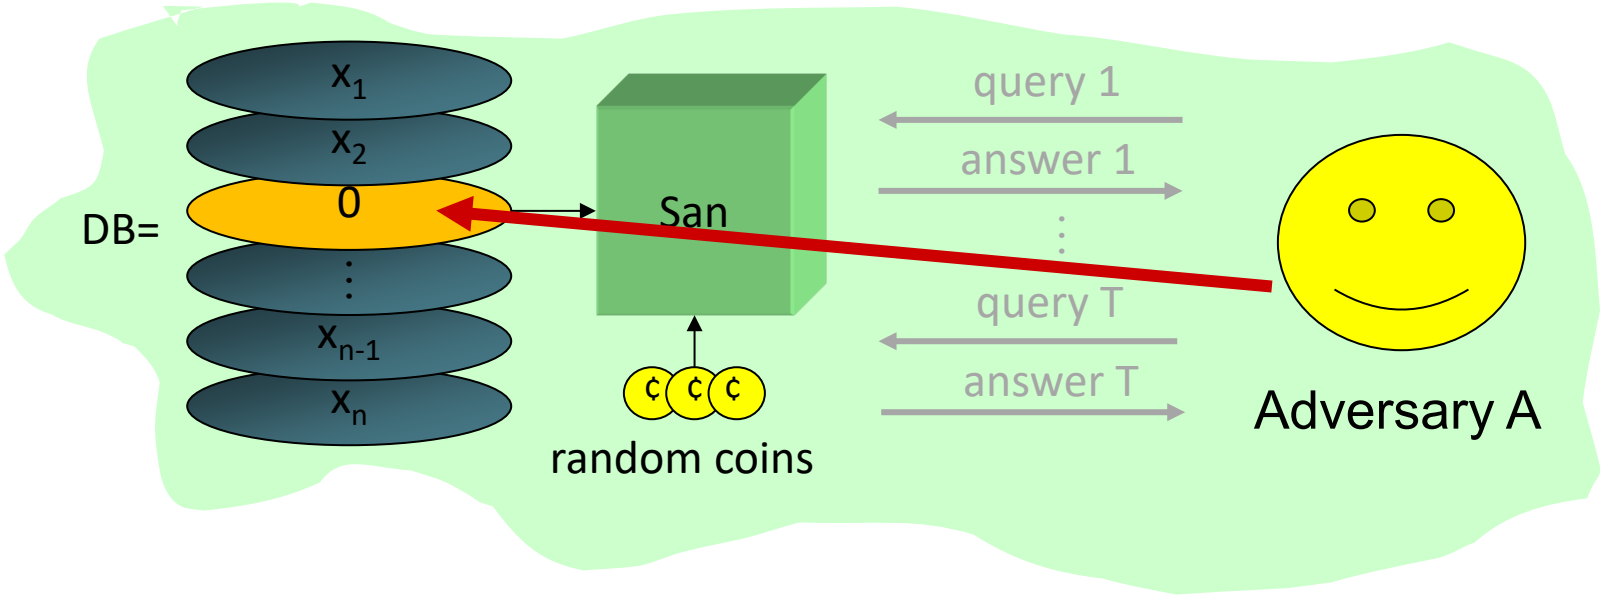


- Define n+1 games
  - Game 0: Adv. interacts with San(DB)
  - Game i: Adv. interacts with San($DB_{-i}$); $DB_{-i} = (x_1,\ldots,x_{i-1},0,x_{i+1},\ldots,x_n)$

# Differential Privacy (2)

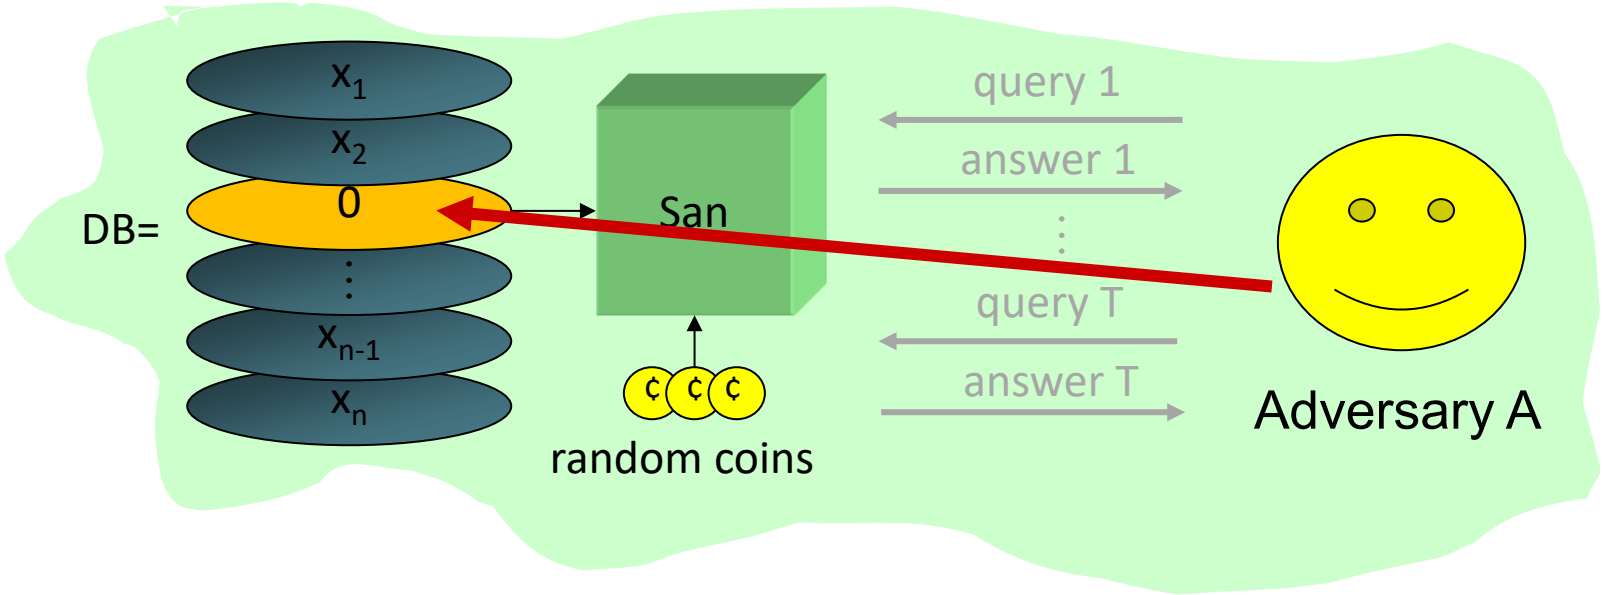


- Define *n*+1 games
  - Game 0: Adv. interacts with San(DB)
  - Game i: Adv. interacts with San($DB_{-i}$); $DB_{-i} = (x_1,\ldots,x_{i-1},0,x_{i+1},\ldots,x_n)$
- Given *S* and prior p() on DB, define *n*+1 posterior distributions

$$p_i(DB|S) = p(DB|S \text{ in Game } i) = \frac{p(San(DB\_i) = S) \times p(DB)}{p(S \text{ in Game } i)}$$

# Differential Privacy (3)

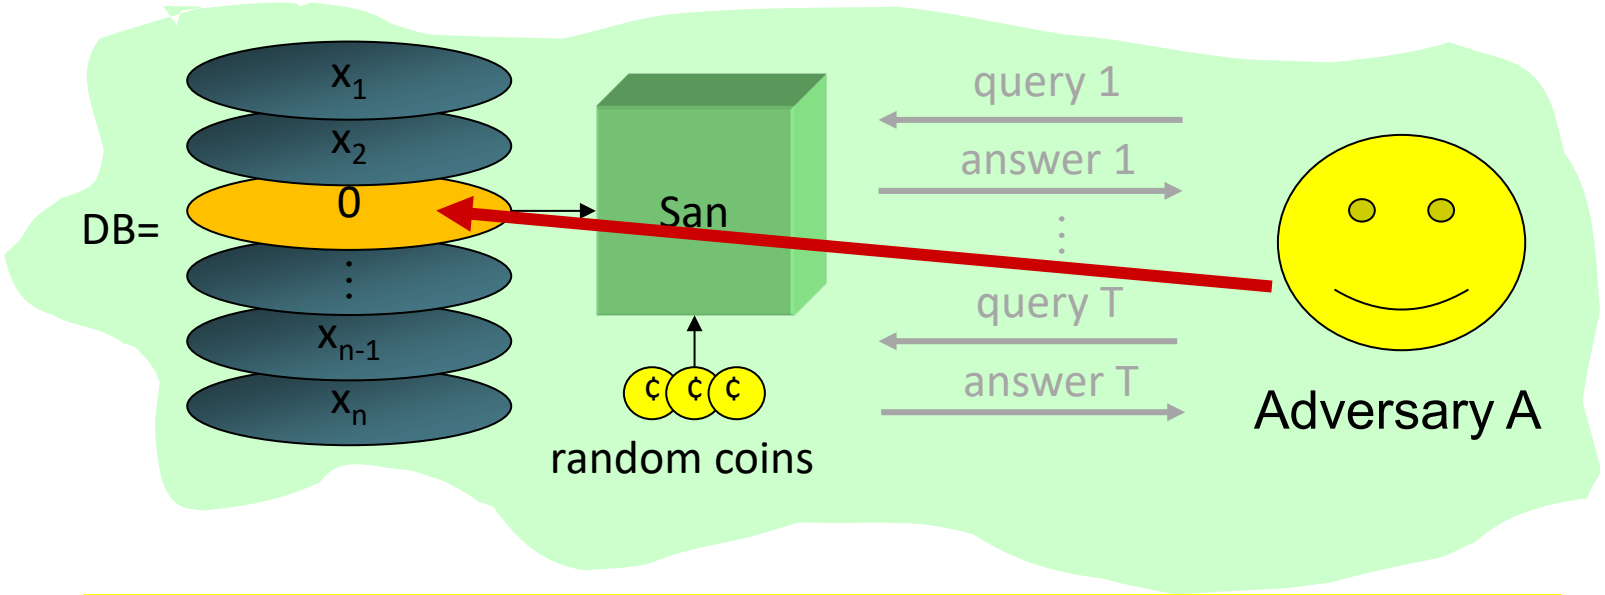


Definition: San is safe if

# Differential Privacy (3)

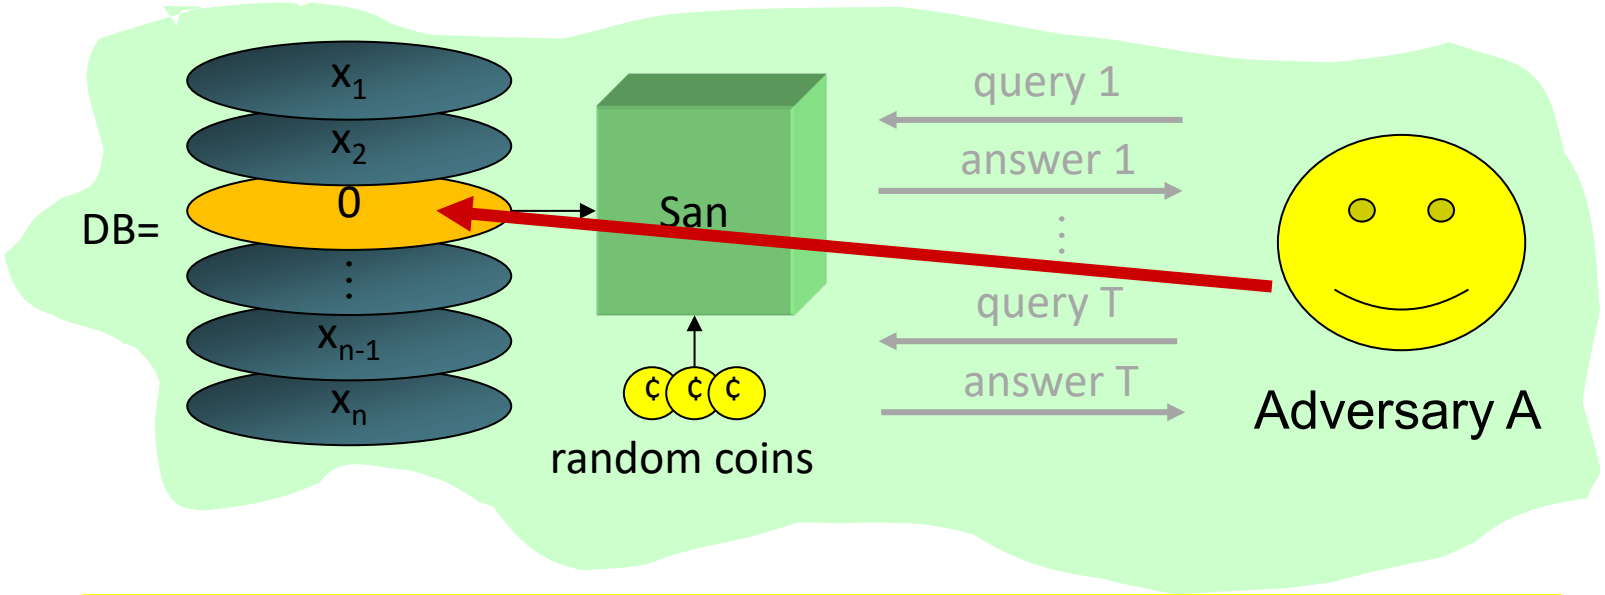


Definition: San is safe if

∀ prior distributions p(¢) on DB,

∀ transcripts *S*, ∀ i =1,…,*n*

# Differential Privacy (3)

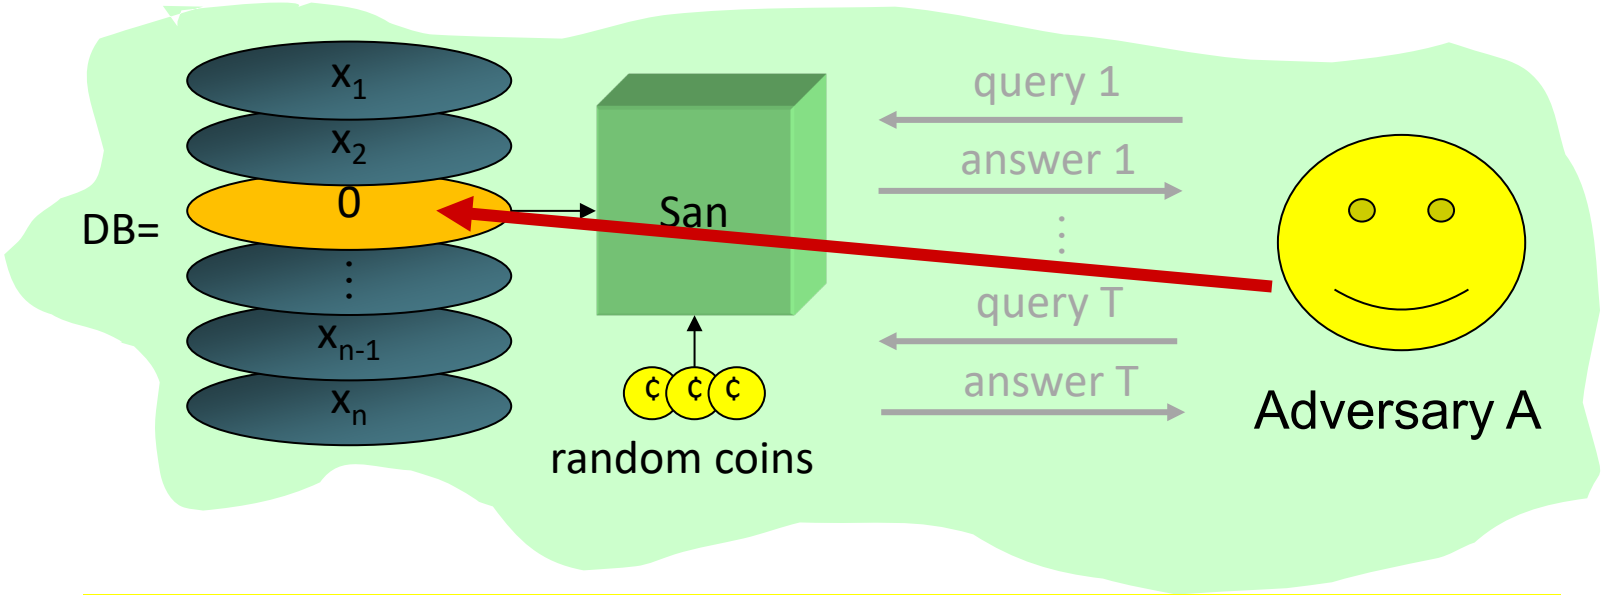


Definition: San is safe if

$\forall$ prior distributions p(¢) on DB,

$\forall$ transcripts *S*, $\forall$ i =1,…,*n*

$$\text{StatDiff}( p_0(¢|S) , p_i(¢|S) ) \leq \varepsilon$$

# Indistinguishability

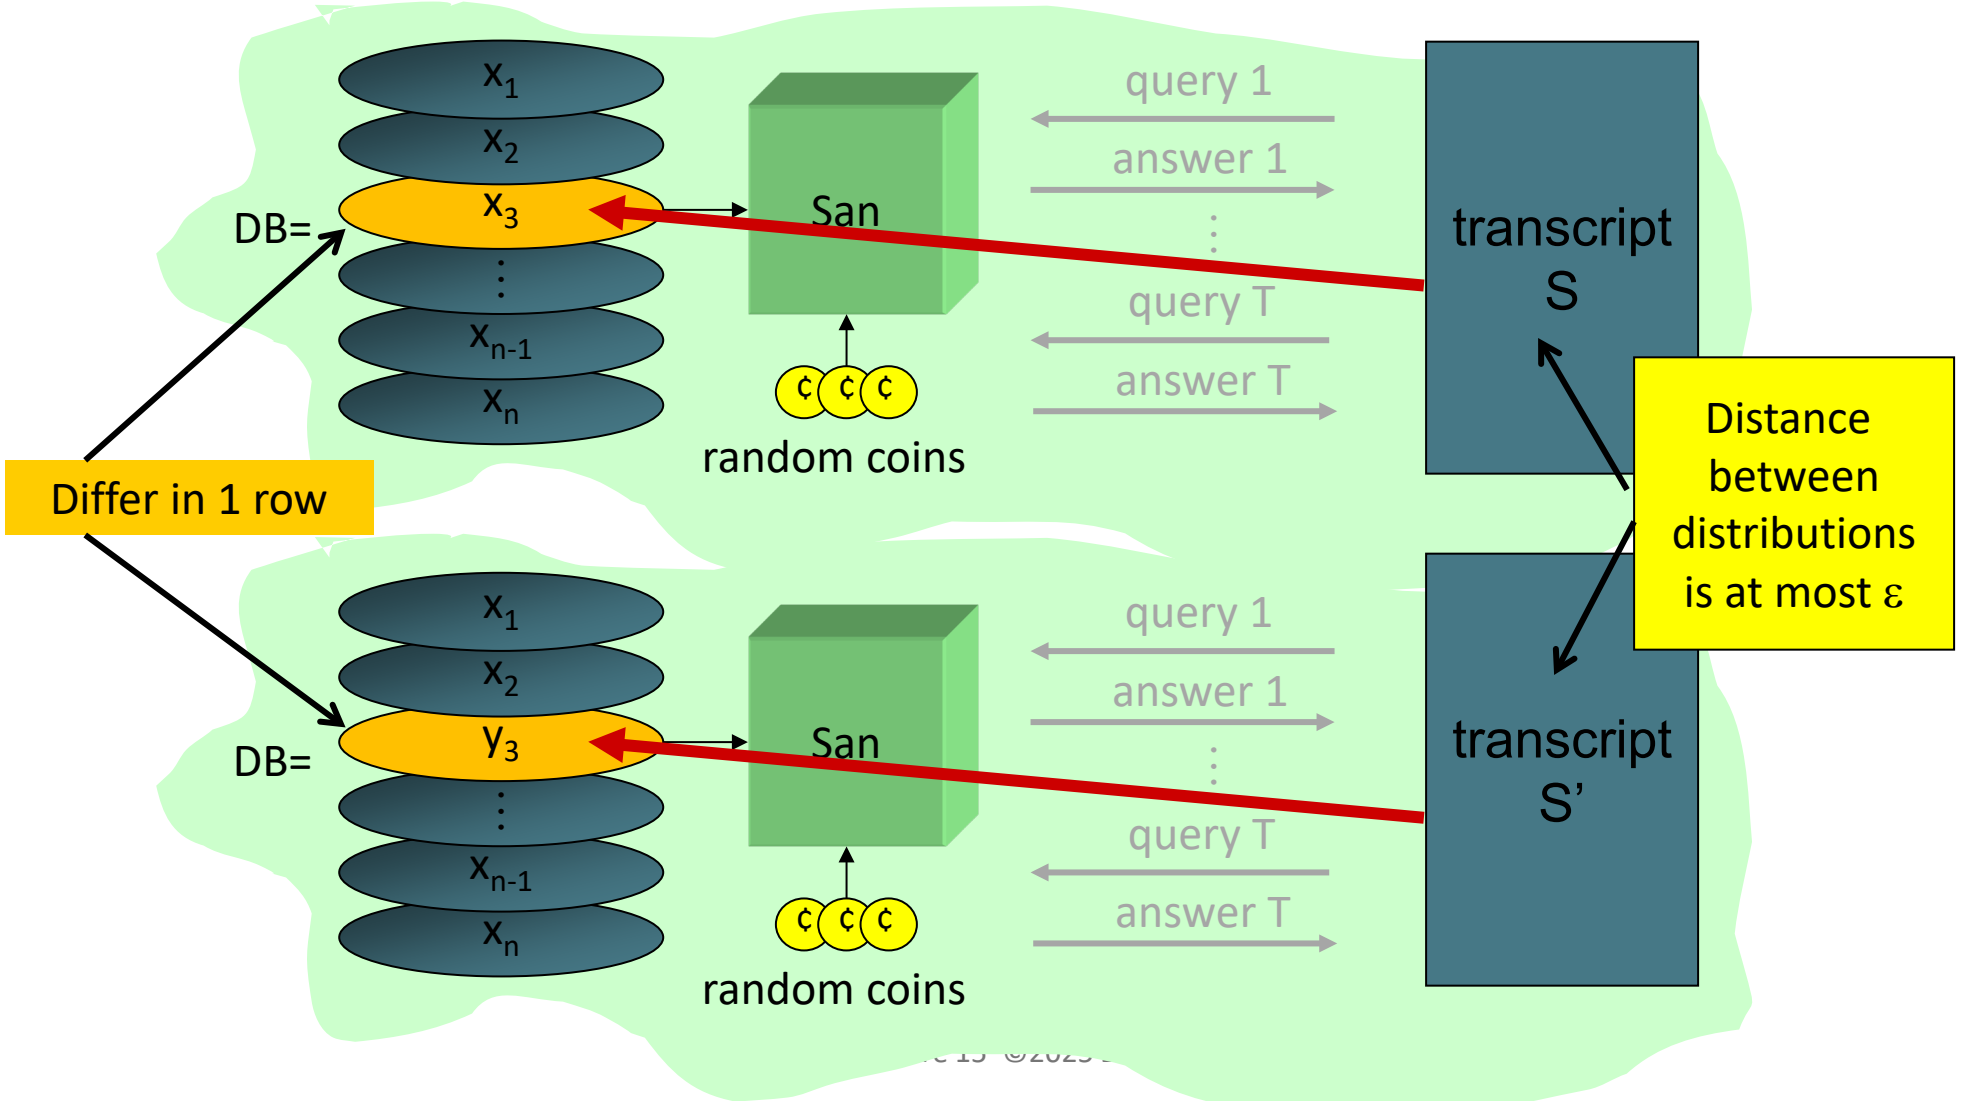

# Which Distance to Use?

- Problem: $\varepsilon$ must be large
  - Any two databases induce transcripts at distance $\leq n\varepsilon$
  - To get utility, need $\varepsilon > 1/n$
- Statistical difference $1/n$ is not meaningful!
- Example: release random point in database
  - $\mathrm{San}(x_1, \ldots, x_n) = (j, x_j)$ for random $j$
- For every $i$ , changing $x_i$ induces statistical difference $1/n$

# (re)Formalizing Indistinguishability


Definition: San is $\varepsilon$-indistinguishable if

$\forall$ A, $\forall$ DB, DB' which differ in 1 row, $\forall$ sets of transcripts S

$$p(\text{San}(DB) \in S) \leq (1 \pm \varepsilon)\, p(\text{San}(DB') \in S)$$

Equivalently, $\forall$ S:

$$\frac{p(\text{San}(DB) = S)}{p(\text{San}(DB') = S)} \leq 1 \pm \varepsilon$$

# Indistinguishability -> Differential Privacy

Definition: San is safe if

$\forall$ prior distributions p(¢) on DB,

$\forall$ transcripts *S*, $\forall$ i =1,…,*n*

$$\text{StatDiff}( p_0(¢|S) , p_i(¢|S) ) \leq \varepsilon$$

# Indistinguishability -> Differential Privacy

Definition: San is safe if

$\forall$ prior distributions p(¢) on DB,

$\forall$ transcripts *S*, $\forall$ i =1,…,*n*

StatDiff( $p_0$(¢|S) , $p_i$(¢|S) ) ≤ $\varepsilon$

$$p_i(DB|S) = p(DB|S \text{ in Game } i) = \frac{p(San(DB_{-i}) = S) \times p(DB)}{p(S \text{ in Game } i)}$$

# Indistinguishability -> Differential Privacy

Definition: San is safe if

∀ prior distributions p(¢) on DB,

∀ transcripts *S*, ∀ i =1,…,*n*

StatDiff( $p_0$(¢|S) , $p_i$(¢|S) ) ≤ ε

$$p_i(DB|S) = p(DB|S \text{ in Game } i) = \frac{p(San(DB_{-i}) = S) \times p(DB)}{p(S \text{ in Game } i)}$$

For every S and DB, indistinguishability implies

# Indistinguishability -> Differential Privacy

Definition: San is safe if

∀ prior distributions p(¢) on DB,

∀ transcripts *S*, ∀ i =1,…,*n*

StatDiff( $p_0(¢|S)$ , $p_i(¢|S)$ ) ≤ ε

$$p_i(DB|S) = p(DB|S \text{ in Game } i) = \frac{p(San(DB_{-i}) = S) \times p(DB)}{p(S \text{ in Game } i)}$$

For every S and DB, indistinguishability implies

$$\frac{p_i(DB|S)}{p_0(DB|S)} = \frac{p(San(DB_{-i}) = S)}{p(San(DB) = S)} \times \frac{p(S \text{ in Game } 0)}{p(S \text{ in Game } i)} \approx 1 \pm 2\epsilon$$

# Indistinguishability -> Differential Privacy

Definition: San is safe if

$\forall$ prior distributions p(¢) on DB,

$\forall$ transcripts *S*, $\forall$ i =1,…,*n*

StatDiff( $p_0$(¢|S) , $p_i$(¢|S) ) ≤ $\varepsilon$

$$p_i(DB|S) = p(DB|S \text{ in Game } i) = \frac{p(San(DB_{-i}) = S) \times p(DB)}{p(S \text{ in Game } i)}$$

For every S and DB, indistinguishability implies

$$\frac{p_i(DB|S)}{p_0(DB|S)} = \frac{p(San(DB_{-i}) = S)}{p(San(DB) = S)} \times \frac{p(S \text{ in Game } 0)}{p(S \text{ in Game } i)} \approx 1 \pm 2\epsilon$$

This implies StatDiff( $p_0$(¢|S) , $p_i$(¢| S) ) ≤ $\varepsilon$

# Differential Privacy in Output Perturbation

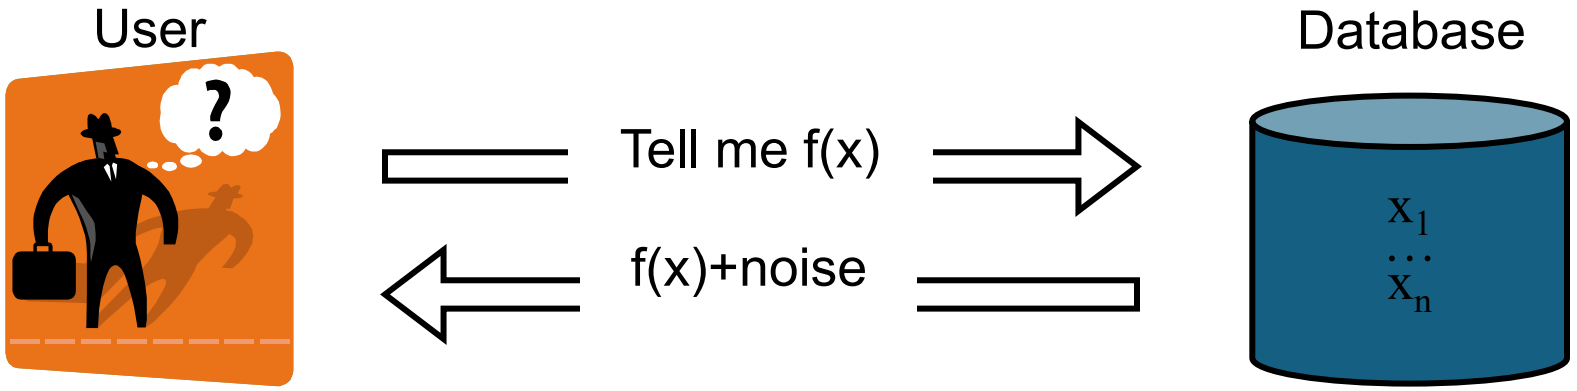


- Intuition: $f(x)$ can be released accurately when $f$ is insensitive to individual entries $x_1, \ldots, x_n$
- Global sensitivity $GS_f = \max_{\text{neighbors } x,x'} \|f(x) - f(x')\|_1$
  - Example: $GS_{average} = 1/n$ for sets of bits
- Theorem: $f(x) + Lap(GS_f / \varepsilon)$ is $\varepsilon$-indistinguishable
  - Noise generated from Laplace distribution

# Sensitivity with Laplace Noise

**Theorem**

*If* $A(x) = f(x) + \mathsf{Lap}\left(\frac{\mathsf{GS}_f}{\varepsilon}\right)$ *then* $A$ *is* $\varepsilon$*-indistinguishable.*

# Sensitivity with Laplace Noise

**Theorem**

*If* $A(x) = f(x) + \mathsf{Lap}\left(\frac{\mathsf{GS}_f}{\varepsilon}\right)$ *then* $A$ *is* $\varepsilon$*-indistinguishable.*

Laplace distribution $\mathsf{Lap}(\lambda)$ has density $h(y) \propto e^{-\frac{\|y\|_1}{\lambda}}$

# Sensitivity with Laplace Noise

**Theorem**

*If* $A(x) = f(x) + \mathsf{Lap}\left(\frac{\mathsf{GS}_f}{\varepsilon}\right)$ *then* $A$ *is* $\varepsilon$*-indistinguishable.*

Laplace distribution $\mathsf{Lap}(\lambda)$ has density $h(y) \propto e^{-\frac{\|y\|_1}{\lambda}}$

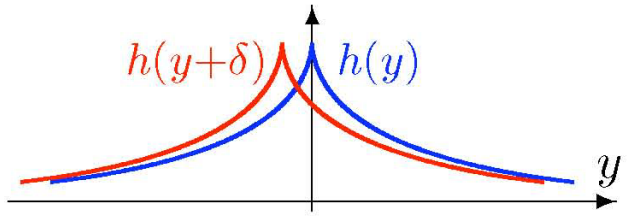

# Sensitivity with Laplace Noise

**Theorem**

*If* $A(x) = f(x) + \mathsf{Lap}\left(\frac{\mathsf{GS}_f}{\varepsilon}\right)$ *then* $A$ *is* $\varepsilon$*-indistinguishable.*

Laplace distribution $\mathsf{Lap}(\lambda)$ has density $h(y) \propto e^{-\frac{\|y\|_1}{\lambda}}$

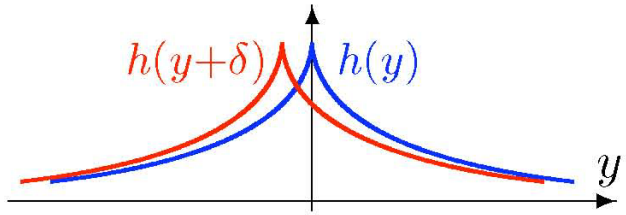


Sliding property of $\mathsf{Lap}\left(\frac{\mathsf{GS}_f}{\varepsilon}\right)$: $\frac{h(y)}{h(y+\delta)} \leq e^{\varepsilon \cdot \frac{\|\delta\|}{\mathsf{GS}_f}}$ for all $y, \delta$

*Proof idea:* $A(x)$: blue curve

# Sensitivity with Laplace Noise

**Theorem**

*If* $A(x) = f(x) + \mathsf{Lap}\left(\frac{\mathsf{GS}_f}{\varepsilon}\right)$ *then* $A$ *is* $\varepsilon$*-indistinguishable.*

Laplace distribution $\mathsf{Lap}(\lambda)$ has density $h(y) \propto e^{-\frac{\|y\|_1}{\lambda}}$

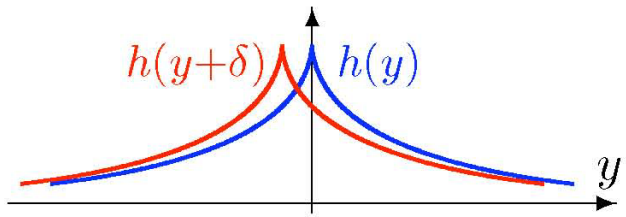


Sliding property of $\mathsf{Lap}\left(\frac{\mathsf{GS}_f}{\varepsilon}\right)$: $\frac{h(y)}{h(y+\delta)} \leq e^{\varepsilon \cdot \frac{\|\delta\|}{\mathsf{GS}_f}}$ for all $y, \delta$

*Proof idea:*

$A(x)$: blue curve

$A(x')$: red curve

# Sensitivity with Laplace Noise

**Theorem**

*If* $A(x) = f(x) + \mathsf{Lap}\left(\frac{\mathsf{GS}_f}{\varepsilon}\right)$ *then* $A$ *is* $\varepsilon$*-indistinguishable.*

Laplace distribution $\mathsf{Lap}(\lambda)$ has density $h(y) \propto e^{-\frac{\|y\|_1}{\lambda}}$

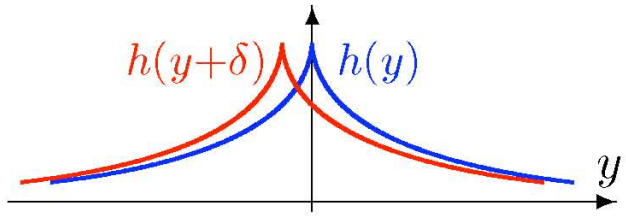


Sliding property of $\mathsf{Lap}\left(\frac{\mathsf{GS}_f}{\varepsilon}\right)$: $\frac{h(y)}{h(y+\delta)} \leq e^{\varepsilon \cdot \frac{\|\delta\|}{\mathsf{GS}_f}}$ for all $y, \delta$

*Proof idea:*

$A(x)$: blue curve

$A(x')$: red curve

$\delta = f(x) - f(x') \leq \mathsf{GS}_f$

# Differential Privacy: Summary

- San gives ε-differential privacy if for all values of DB and Me and all transcripts $t$:

$$\frac{\Pr(\mathit{San}(\mathrm{DB} - \mathrm{Me}) = t)}{\Pr(\mathit{San}(\mathrm{DB} + \mathrm{Me}) = t)} \leq e^{\varepsilon} \approx 1 \pm \varepsilon$$

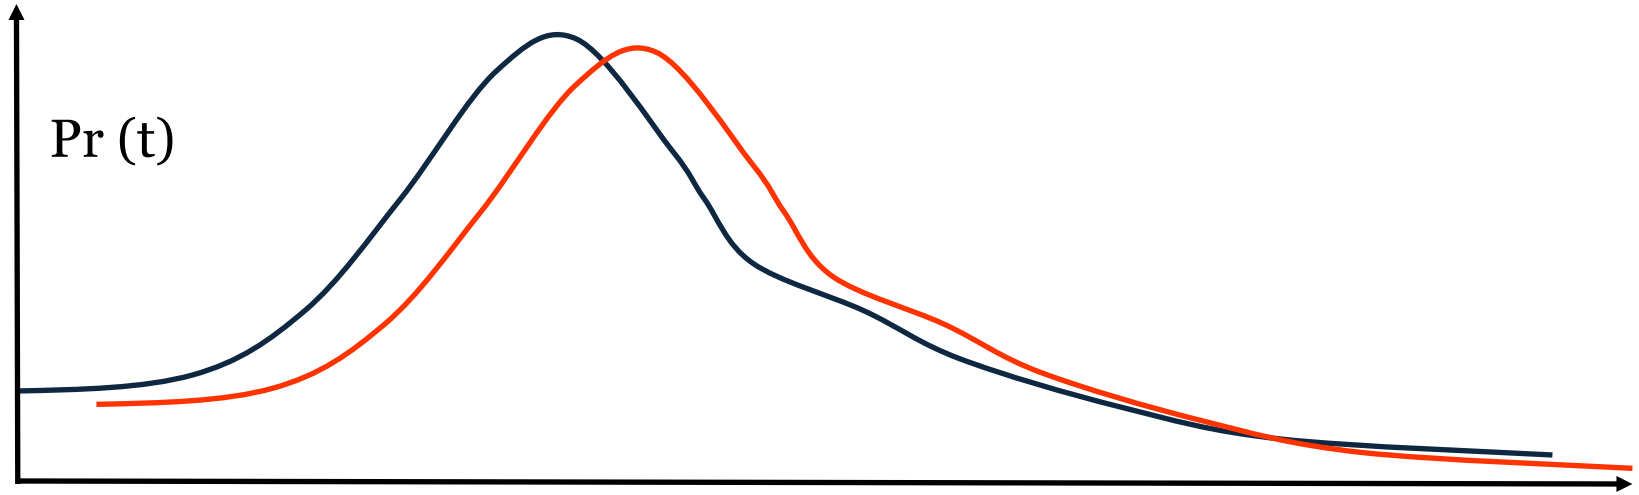

# Differential Privacy

***No perceptible risk is incurred by joining DB***
***Anything adversary can do to me, it could do without me (my data)***

Neutralizes all linkage attacks.
Composes unconditionally and automatically: $\Sigma_i \varepsilon_i$

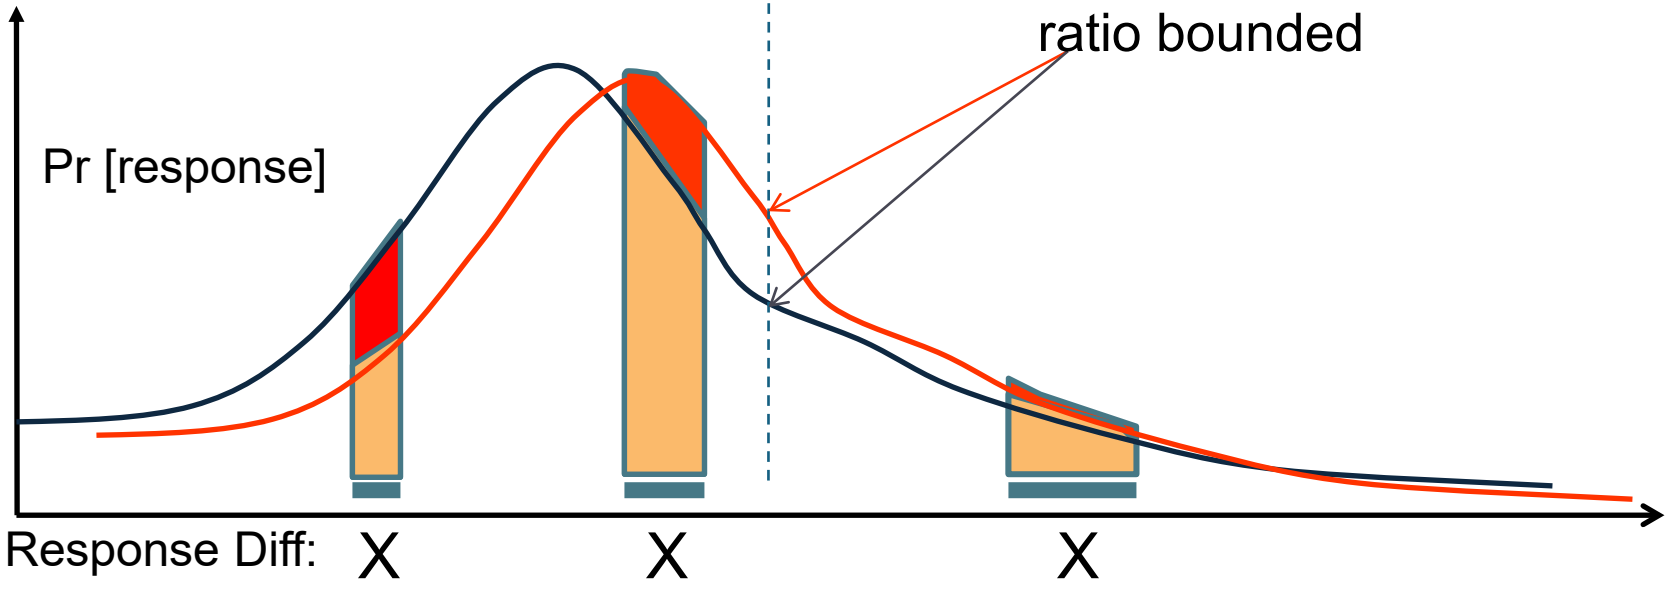

# Access Control

- Security Rule
- Pillars of Security
- Variations of Access Control

# HIPAA – Data Protection

PRIVACY RULE (2002)

Dept of Health & Human Services. Standards for privacy of individually identifiable health information; Final Rule. Federal Register. 45 CFR: Pt 160 and 164.

SECURITY RULE (2003)

Dept of Health & Human Services. Standards for the Protection of Electronic Health Information; Final Rule. Federal Register. 45 CFR: Pt 164.

# HIPAA Security Rule

- Administrative Safeguards
- Physical Safeguards
- **Technical Safeguards**
- Organizational Requirements

http://www.cms.hhs.gov/SecurityStandard/

# Administrative Safeguards

| Standards | Implementation Specification | Required vs. Addressable |
|---|---|---|
| **Security Management Process** | **Risk Analysis** | **R** |
| | **Risk Management** | **R** |
| | **Sanction Policy** | **R** |
| | **Information System Activity Review** | **R** |
| **Assigned Security Responsibility** | | **R** |
| Workforce Security | Authorization and/or Supervision | A |
| | Workforce Clearance Procedure | A |
| | Termination Procedures | A |
| **Information Access Management** | **Isolating Healthcare Clearinghouse Function** | **R** |
| | Access Authorization | A |
| | Access Establishment and Modification | A |
| Security Awareness & Training | Security Reminders | A |
| | Protection from Malicious Software | A |
| | Log-in Monitoring | A |
| **Security Incident Procedures** | **Response and Reporting** | **R** |
| Contingency Plan | Data Backup Plan | R |
| | Disaster Recovery Plan | R |
| | Emergency Mode Operation Plan | R |

# Physical Safeguards

| Standards | Implementation Specification | Required vs. Addressable |
|---|---|---|
| Facility Access Controls | Contingency Operations | A |
| | Facility Security Plan | A |
| | Access Control and Validation Procedures | A |
| | Maintenance Records | A |
| **Workstation Use** | | **R** |
| **Workstation Security** | | **R** |
| **Device & Media Controls** | **Disposal** | **R** |
| | **Media Reuse** | **R** |
| | Accountability | A |
| | Data Backup & Storage | A |

# Technical Safeguards

| Standards | Implementation Specification | Required vs. Addressable |
|---|---|---|
| **Access Control** | **Unique User Identification** | **R** |
| | **Emergency Access Procedure** | **R** |
| | Automatic Logoff | A |
| | Encryption and Decryption | A |
| **Audit Controls** | | **R** |
| Integrity | Mechanism to Authenticate ePHI | A |
| **Person or Entity Authentication** | | **R** |
| Transmission Security | Integrity Controls | A |
| | Encryption | A |

# Three Pillars of Security

Least Privilege

Separation of Duties

Data Abstraction

# Least Privilege

User should be provided with no more privileges than are necessary to perform their job

# Separation of Duties

Requirement for multiple types of individuals to complete a task

# Data Abstraction

Permissions are related to the type of data being handled

# Access Control – What to Control?

- Subjects S (or Users)
- Objects *O* (or Patients)
- Rights *R*
  - ☐ Read from Record
  - ☐ Issue Order
  - ☐ Write to Record
  - ☐ Request Consult
  - ☐ Could specialize to “type” of information
    - demographics
    - diagnoses
    - treatments

# Subjects & Principals

- One-to-many mapping of subjects to principals
- Intention is to ensure *accountability* for one's actions

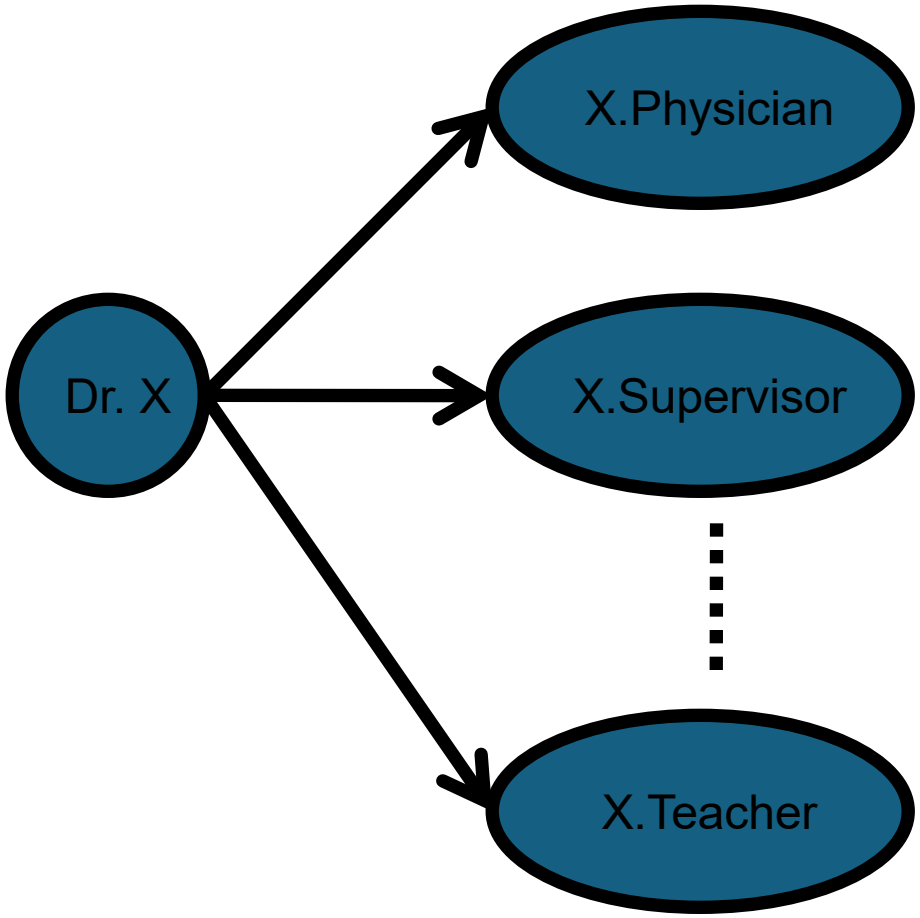

# Many Variations

- **Access Matrix (AM)**
- Mandatory Access Control (MAC)
- Discretionary Access Control (DAC)
- Role-Based Access Control (RBAC)

# Abstraction of Access Control
(Lampson 1971)

- A right is a relation for subjects and objects

$$r(s,o) \subseteq \mathrm{Rel}$$

- Specification of which rights can be invoked by which subject for which object

B. Lampson. Protection. Proc. 5th Princeton Conference on Information Sciences and Systems. 1971: 437.

# Access Matrix

| | **Object** | | |
|---|---|---|---|
| **Subject** | A | B | C |
| Dr. D | *R-X* | *RWX* | *RWX* |
| Nurse E | *R-X* | *RWX* | *R-X* |
| Biller F | *R--* | *R--* | *R--* |

r(Dr. D, C) =
{R, W, X}

# How to Use an Access Matrix

- Can allow for dynamic protections
  - Operations for assignment & revocation of rights

- Can permit “special” rights:
  - Ownership of object
  - Copy of object
  - Control of rights modification for object
  - …

# Views on the Matrix

- Access Control List (ACL)
  - For a single object
  - Indicates which subject can invoke which right

| Subject | Object A |
|---|---|
| Dr. D | $R, X$ |
| Nurse E | $R, X$ |
| Biller F | $R$ |

# Views on the Matrix

- Capability List
  - For a single subject
  - Indicates which rights can be invoked by the subject across objects which right

| | **Object** | | |
|---|---|---|---|
| **Subject** | A | B | C |
| Dr. D | $R, X$ | $R, W, X$ | $R, W, X$ |

# Many Variations

- Access Matrix (AM)
- **Mandatory Access Control (MAC)**
- **Discretionary Access Control (DAC)**
- Role-Based Access Control (RBAC)

# Mandatory vs. Discretionary

- Mandatory access controls (MAC) restrict the access of subjects to objects on the basis of “security” labels

- Discretionary access controls (DAC) permits access rights to be propagated from one subject to another
  - Possession of an access right by a subject is sufficient to allow access to the object

# Take a Step Back

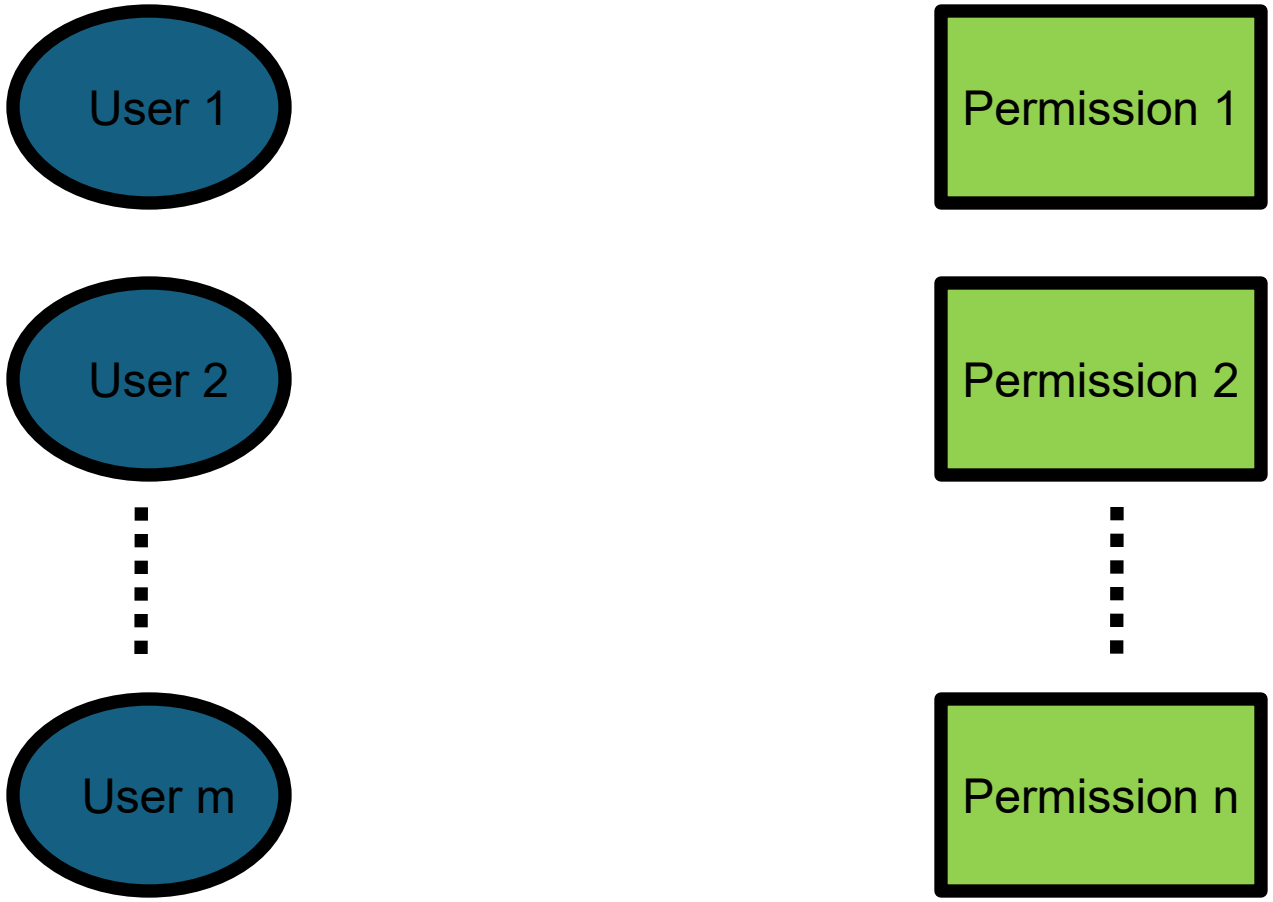

# Rights (or Permission) Assignment

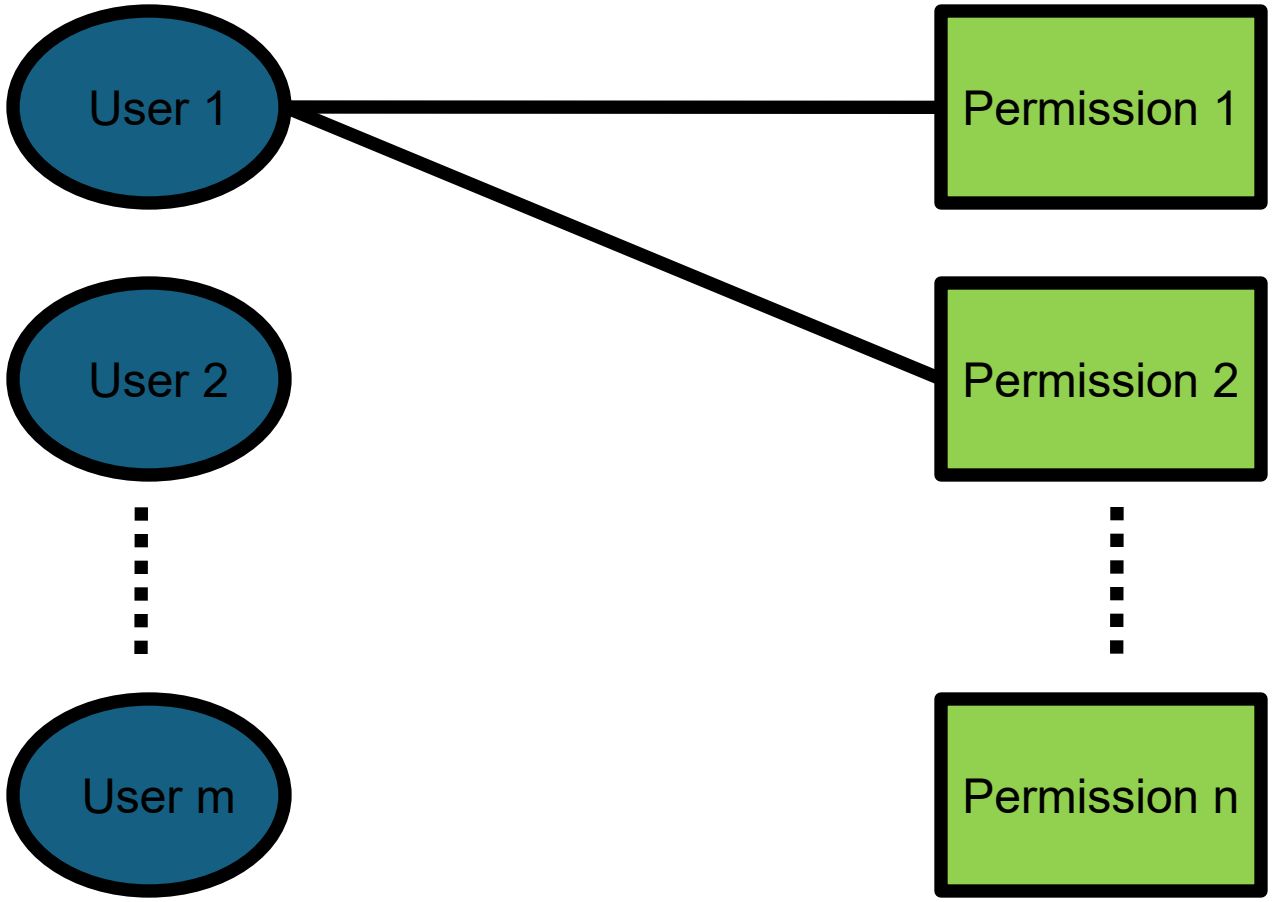

# Rights (or Permission) Assignment

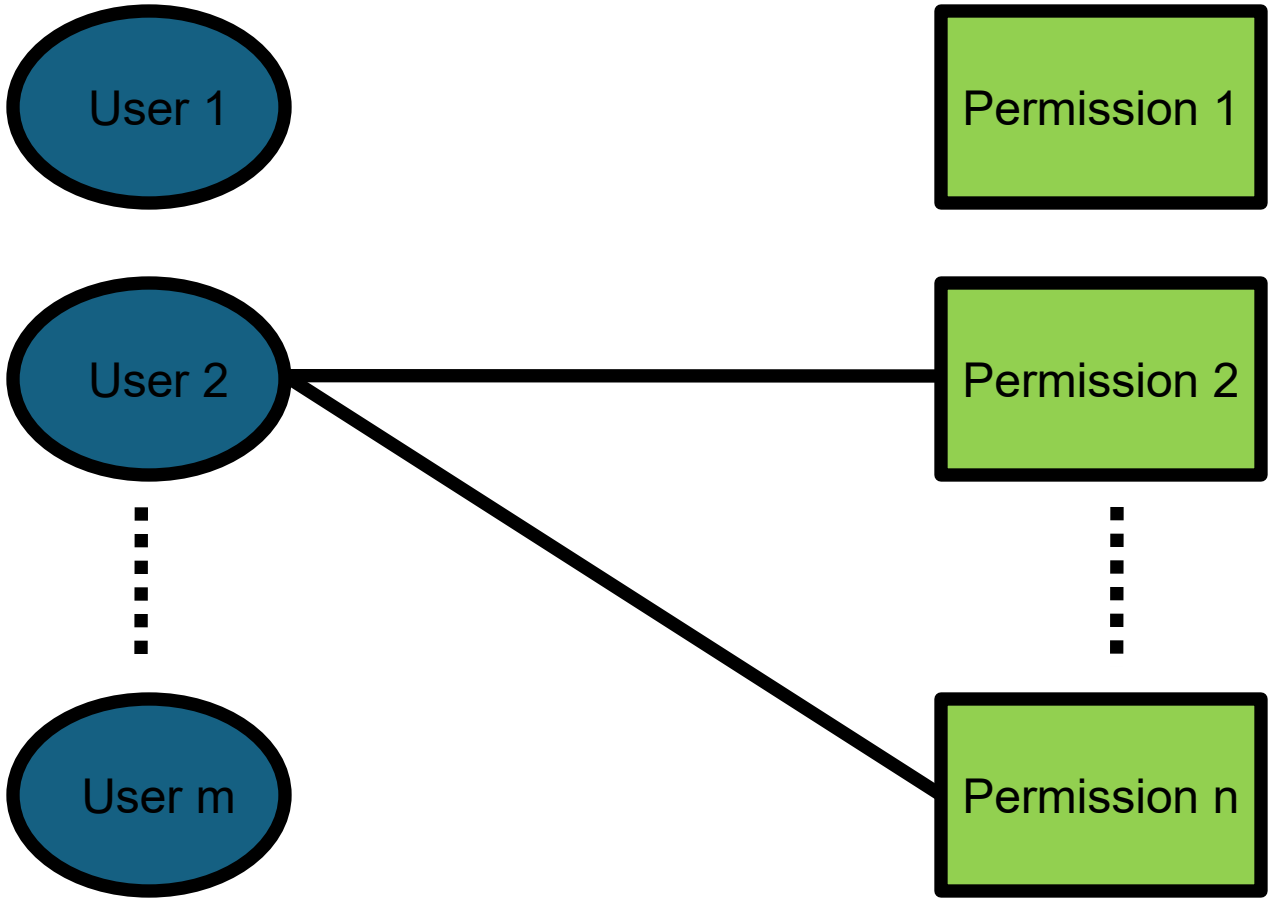

# Delegation

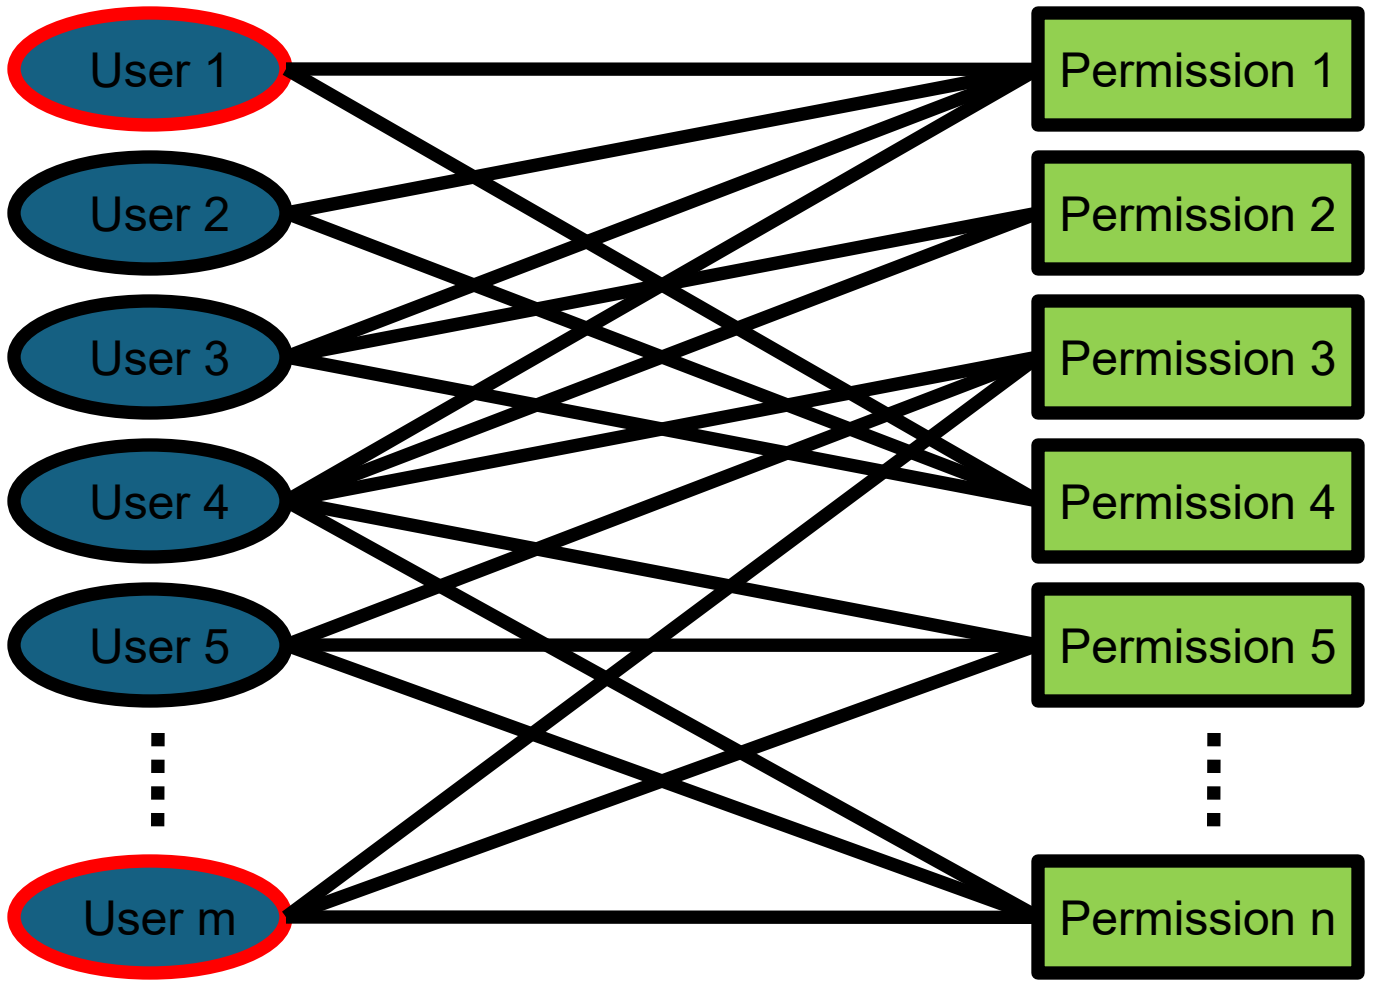
User 1
User 2
User 3
User 4
User 5
User m
Permission 1
Permission 2
Permission 3
Permission 4
Permission 5
Permission n

# Delegation

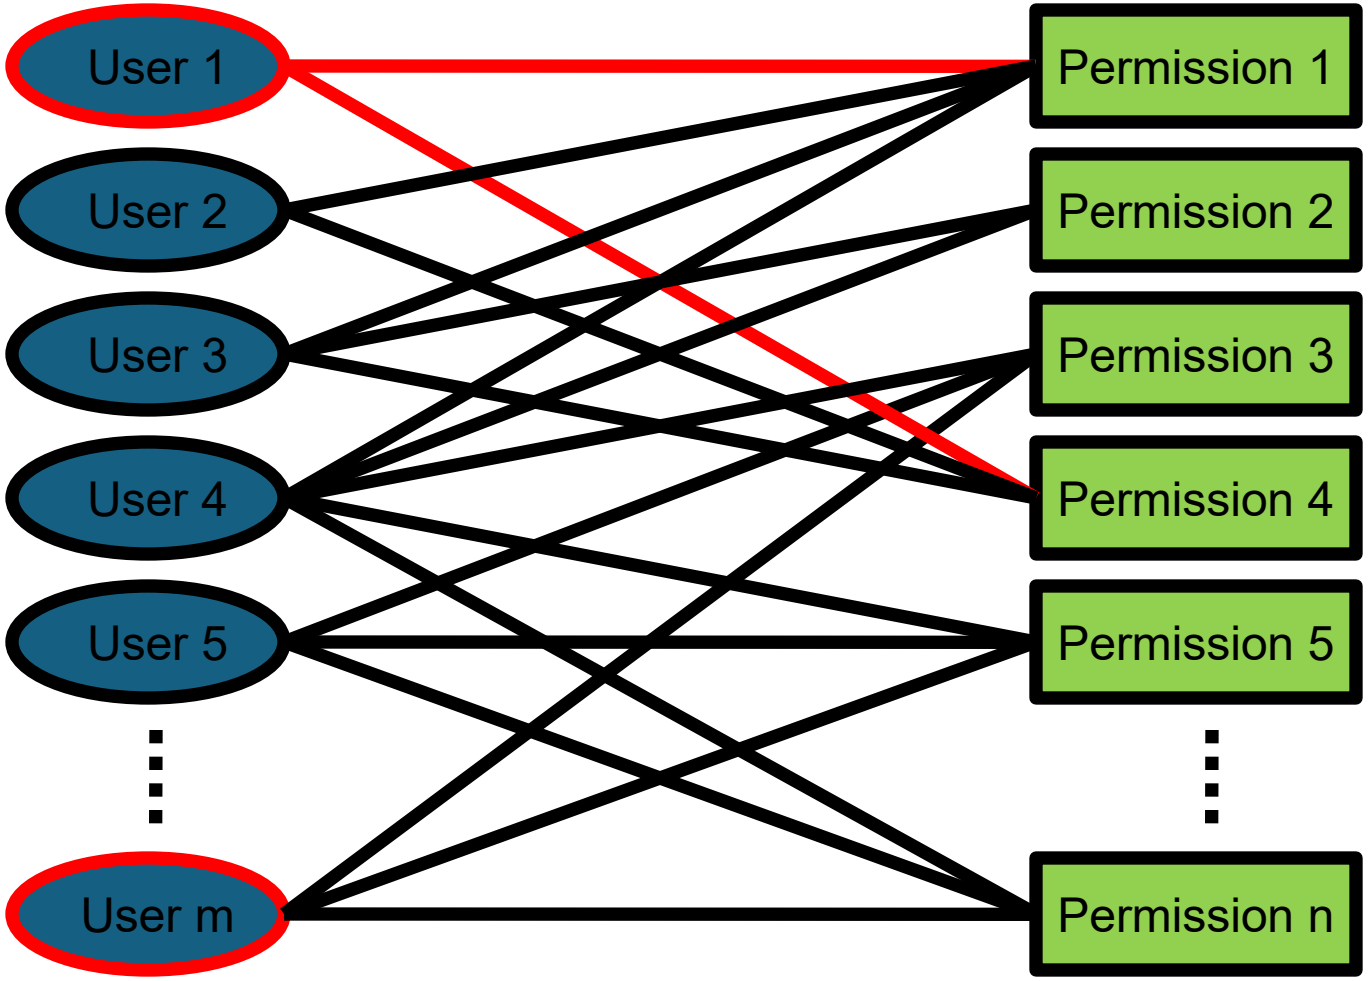
User 1
User 2
User 3
User 4
User 5
User m
Permission 1
Permission 2
Permission 3
Permission 4
Permission 5
Permission n

# Delegation

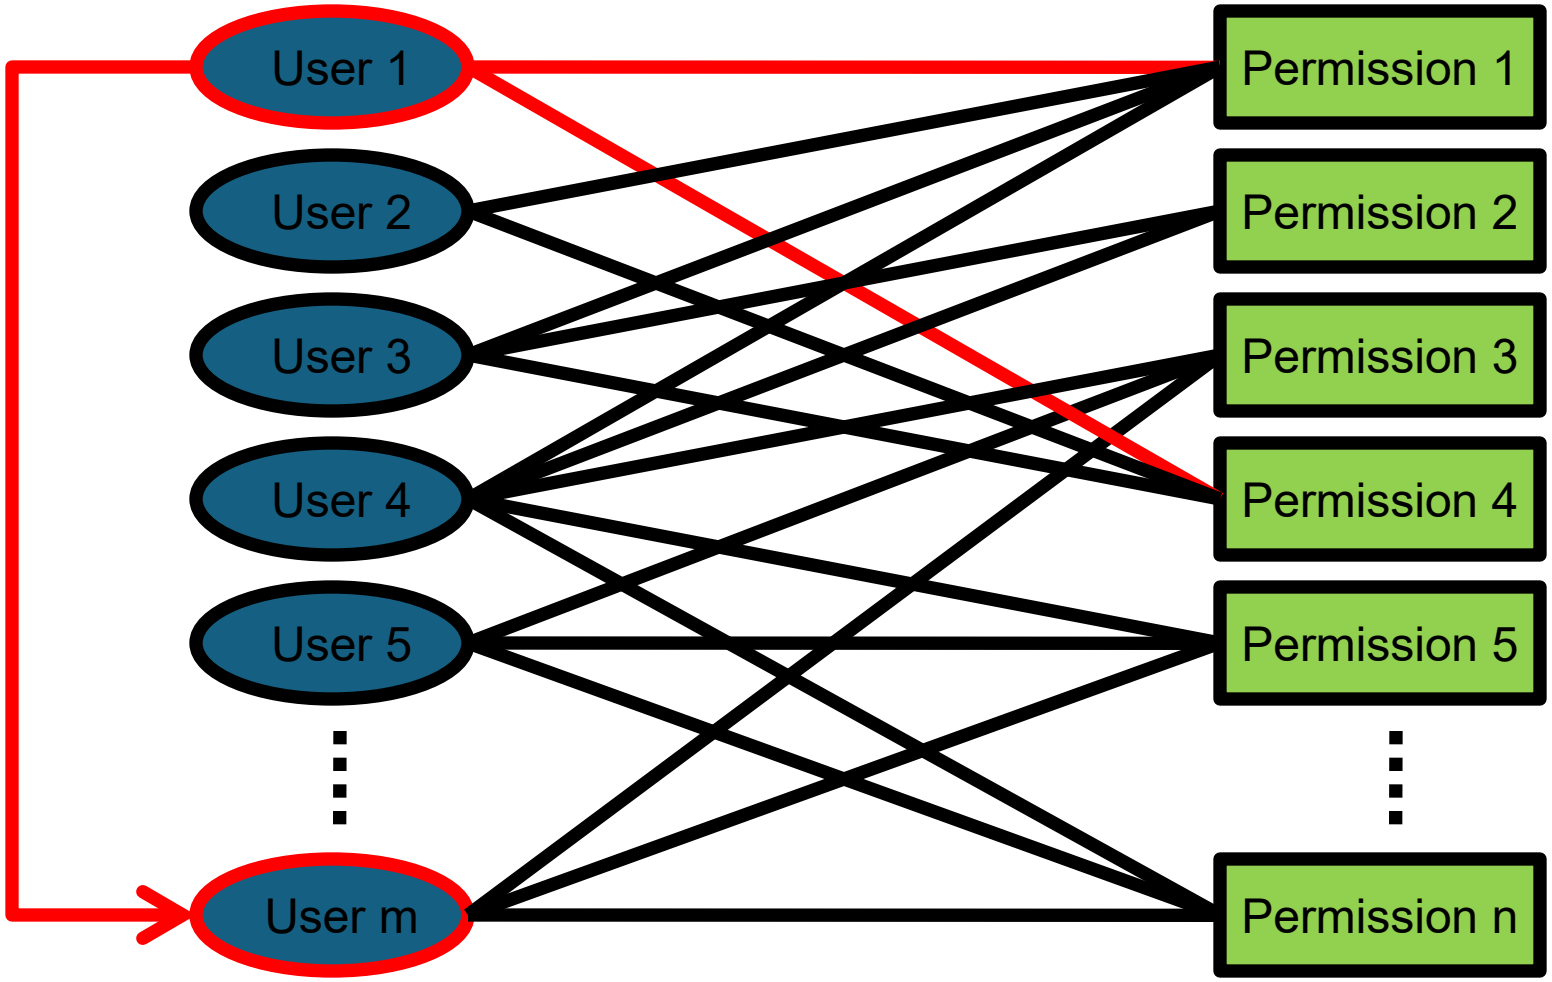
User 1
User 2
User 3
User 4
User 5
User m
Permission 1
Permission 2
Permission 3
Permission 4
Permission 5
Permission n

# User-Permission Relation via Transitivity

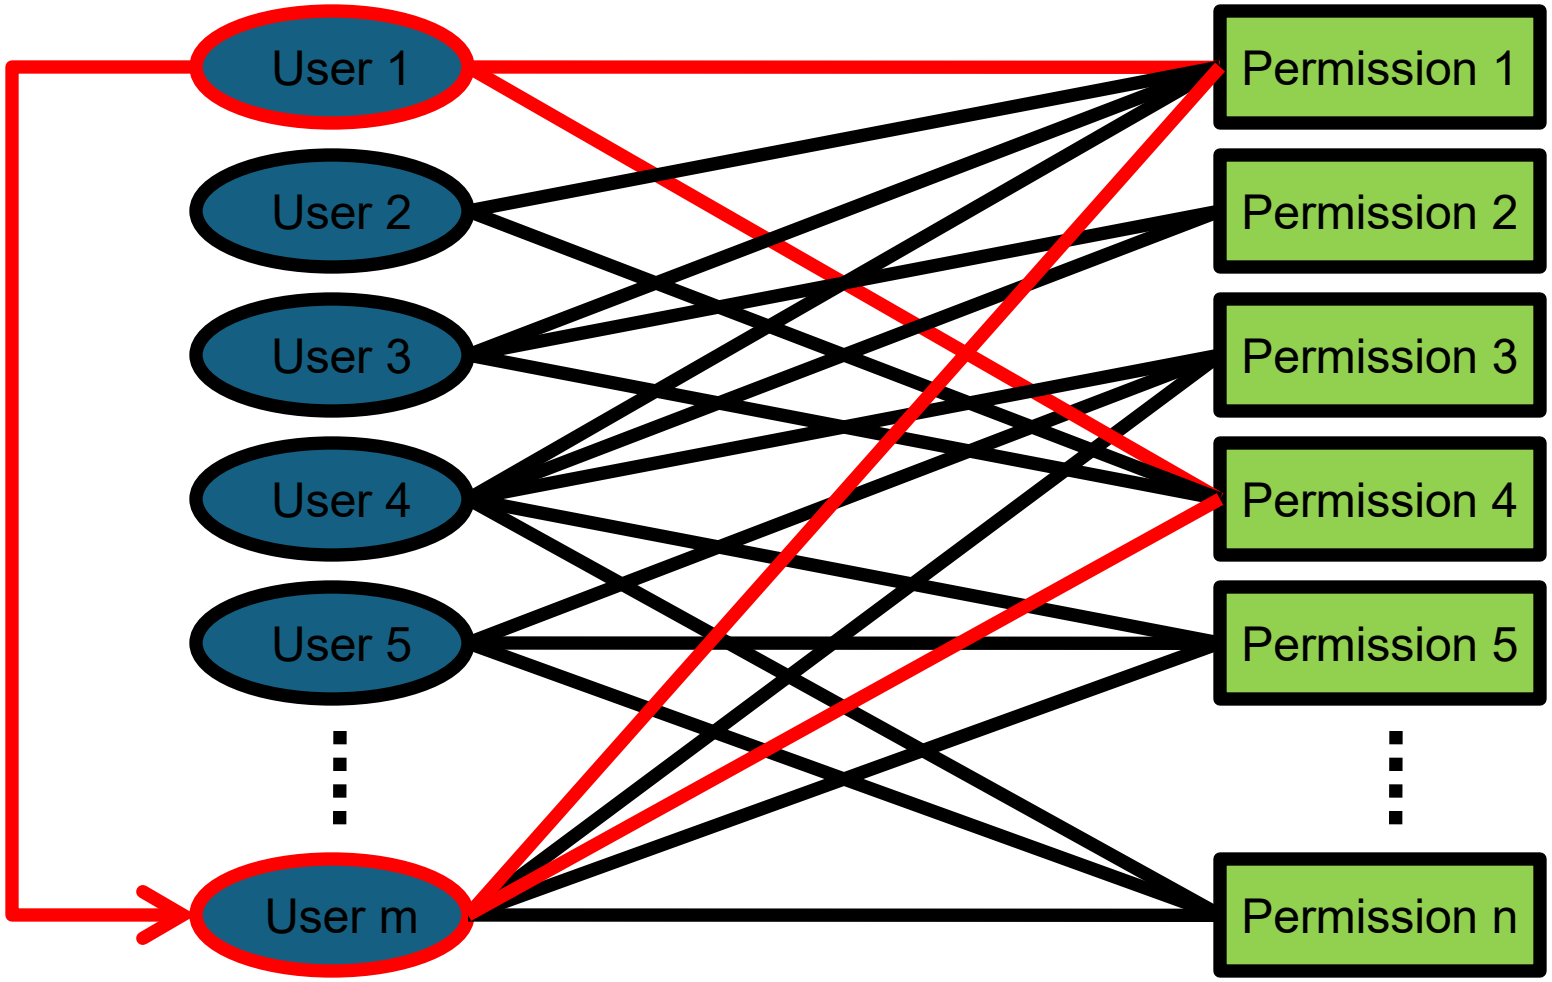

# Many Variations

- Access Matrix (AM)
- Mandatory Access Control (MAC)
- Discretionary Access Control (DAC)
- **Role-Based Access Control (RBAC)**

# Many Potential Assignments

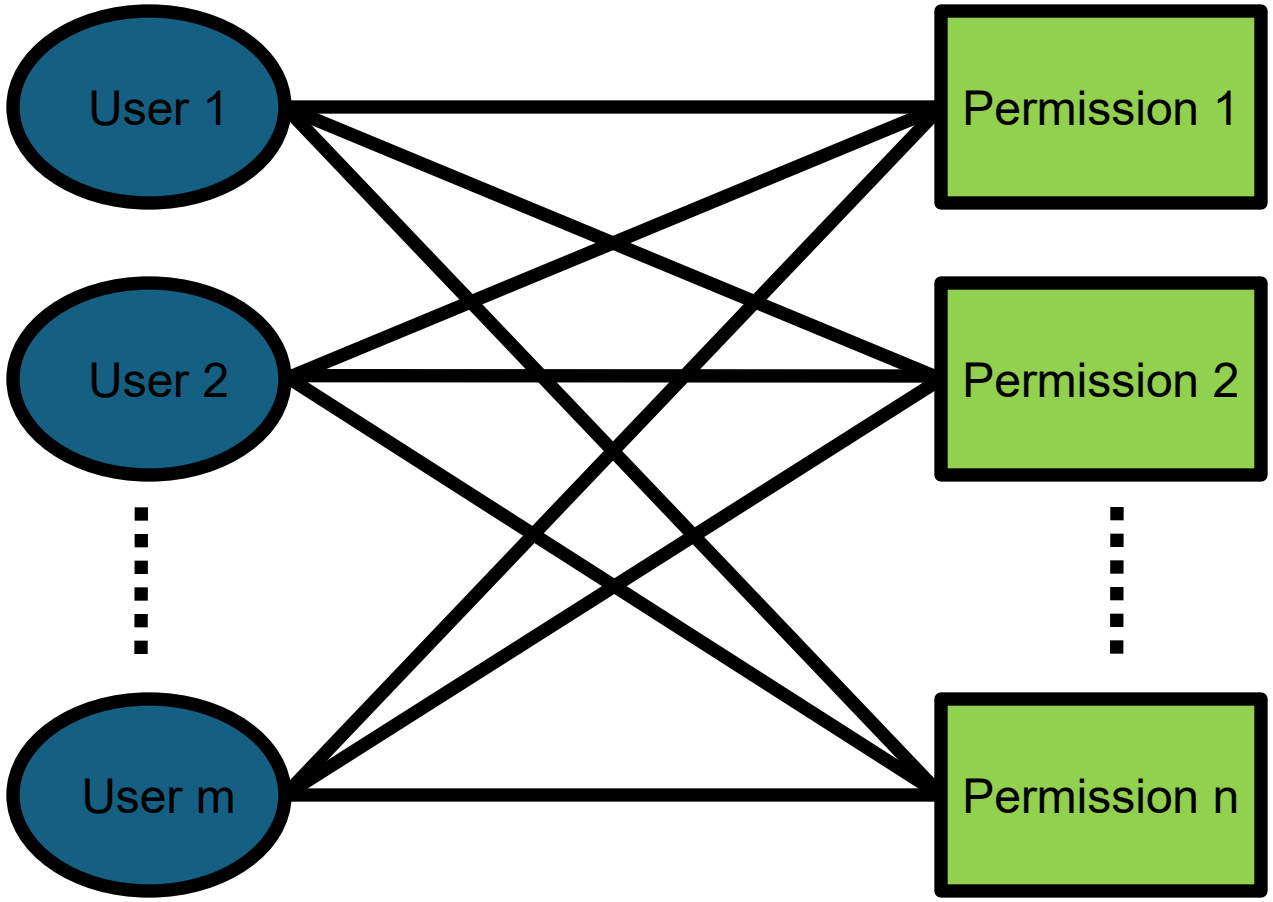

# Role-Based Access Control (RBAC)

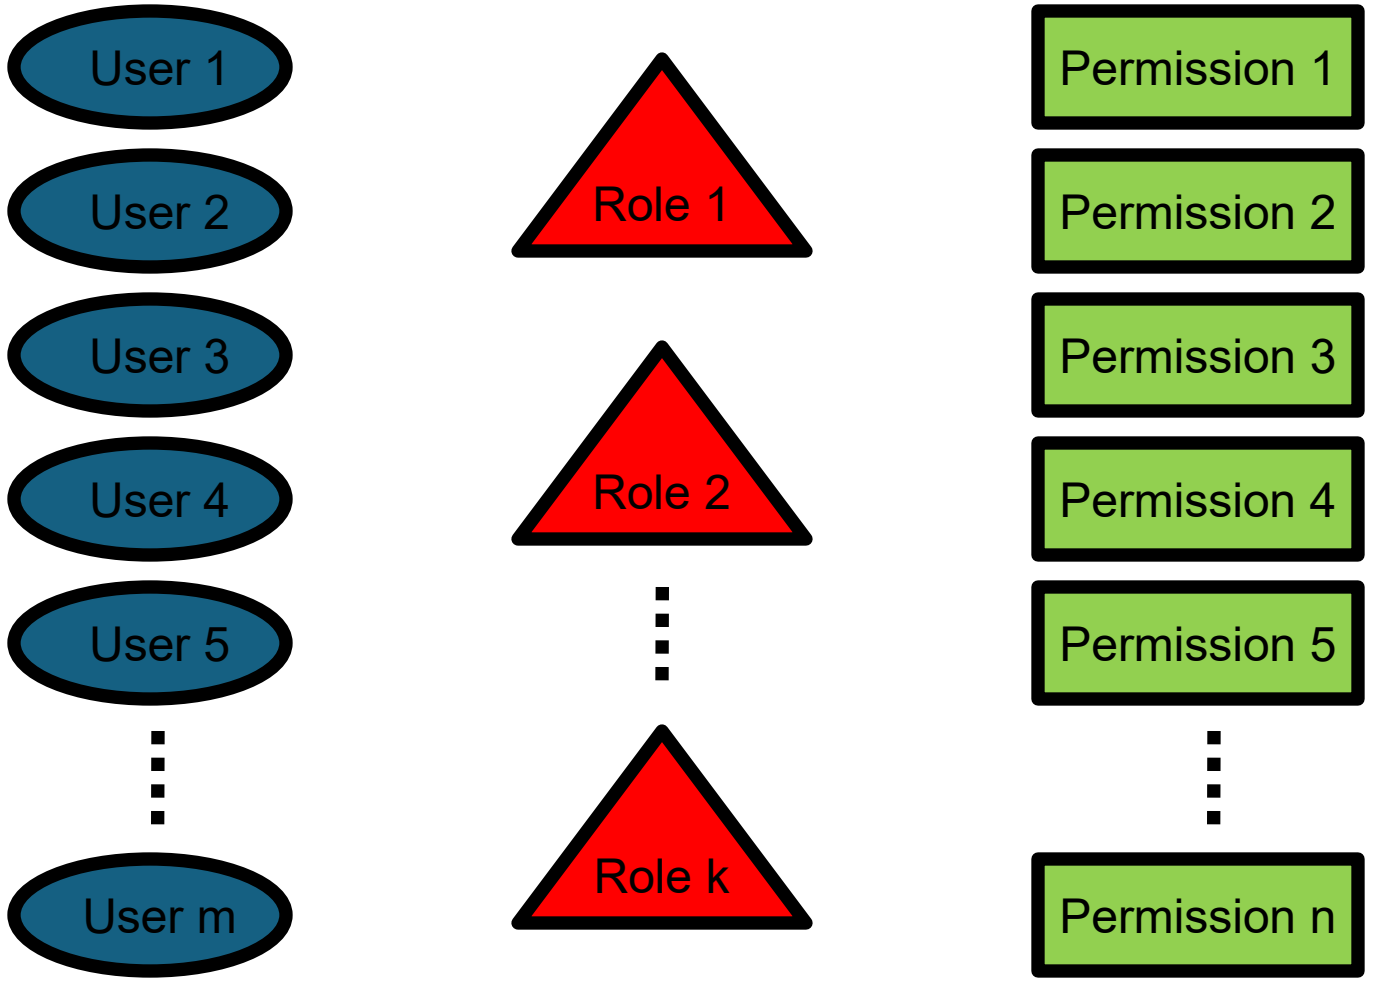

# Role-Based Access Control (RBAC)

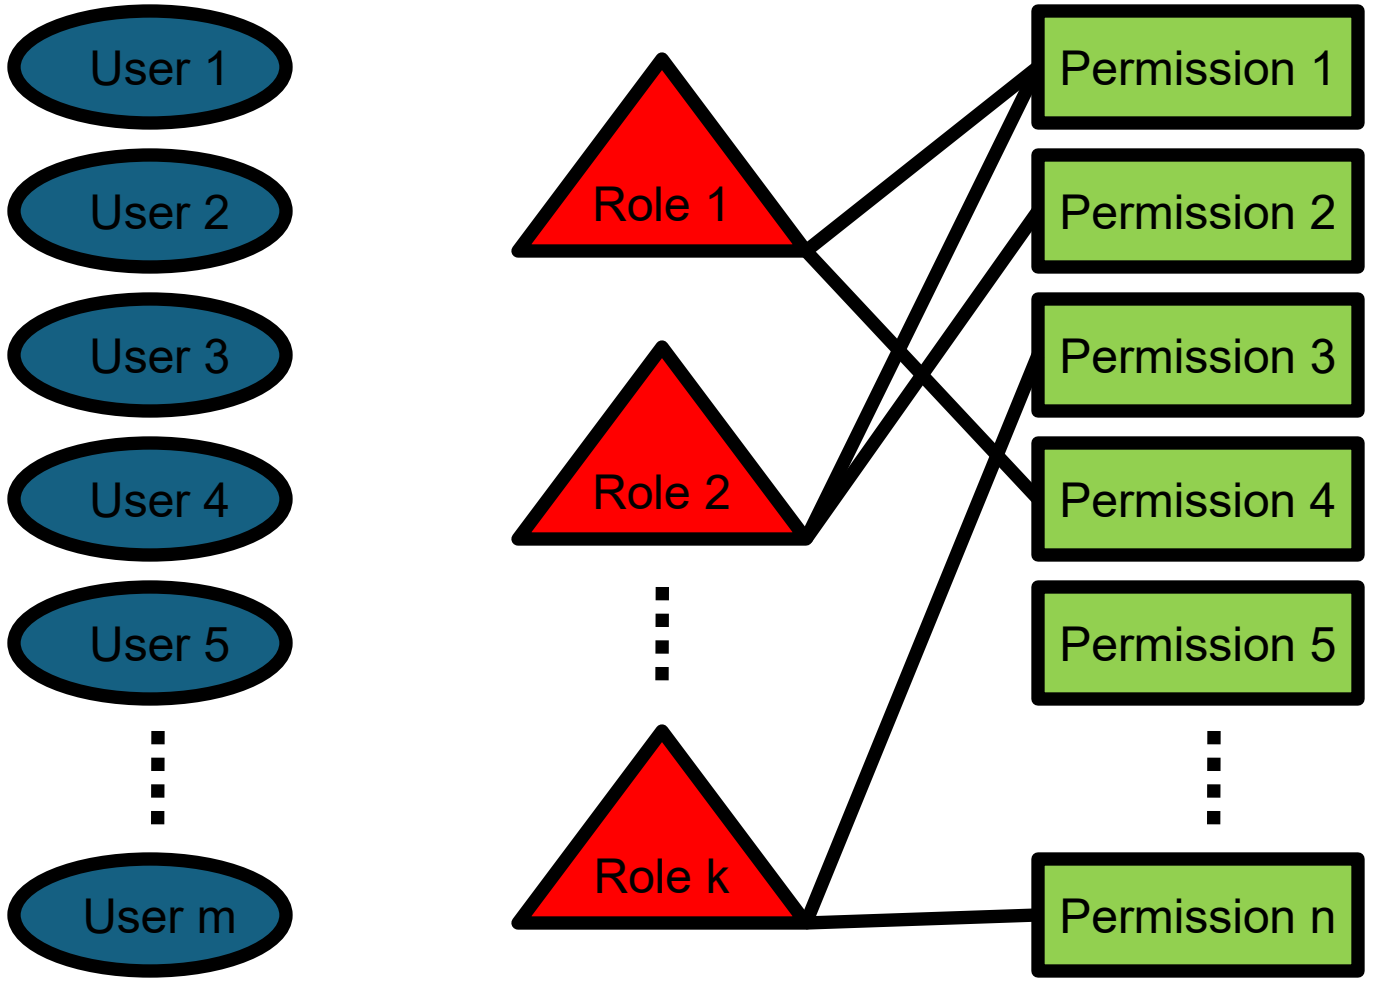

# Can Map Users to Roles

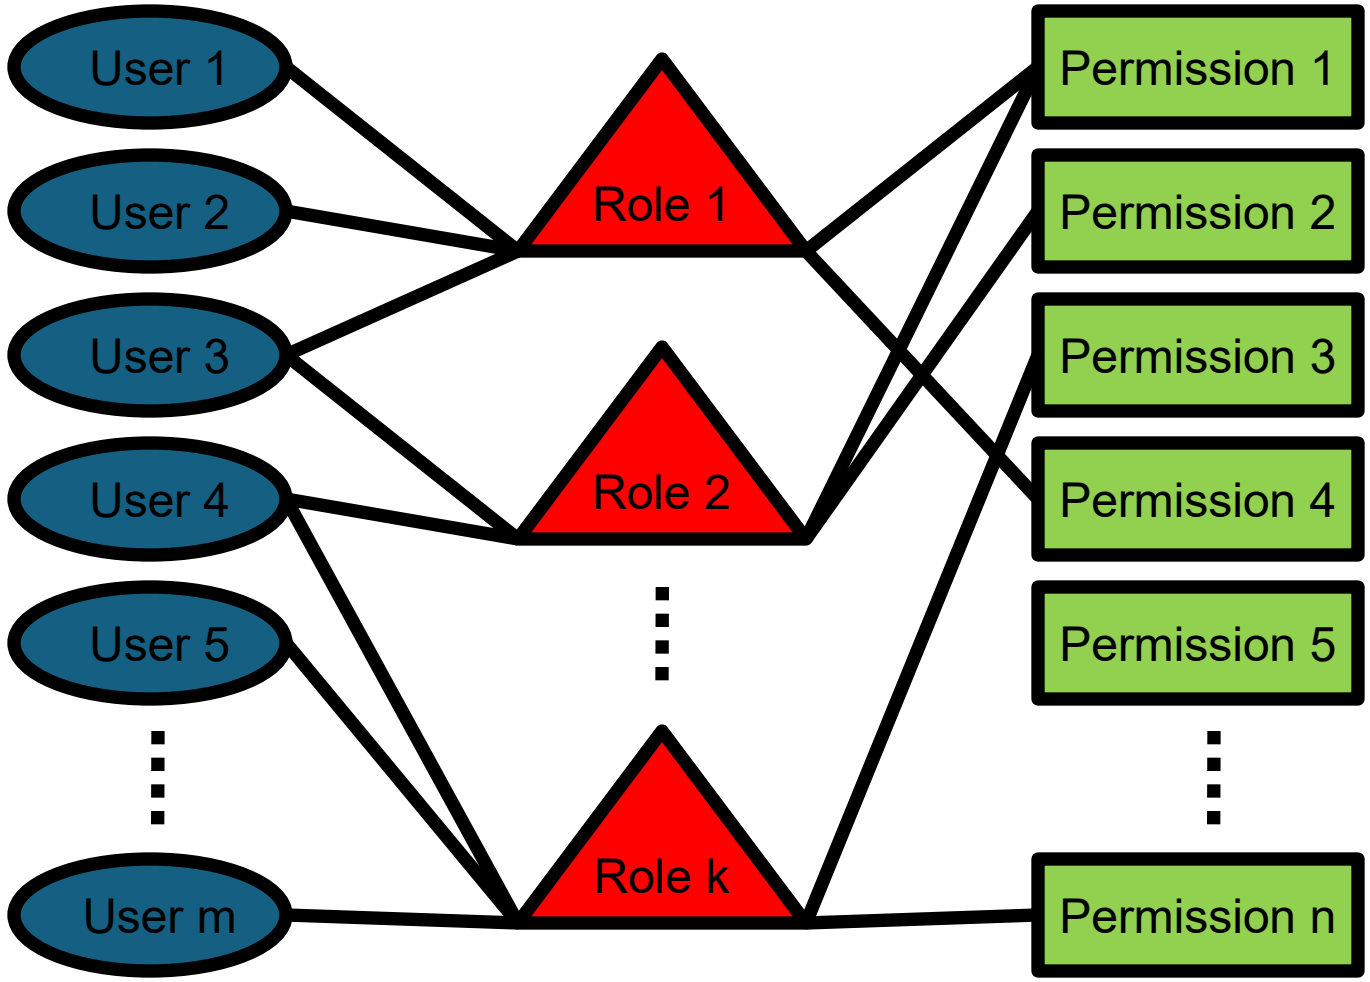

# Users can Have Multiple Roles!

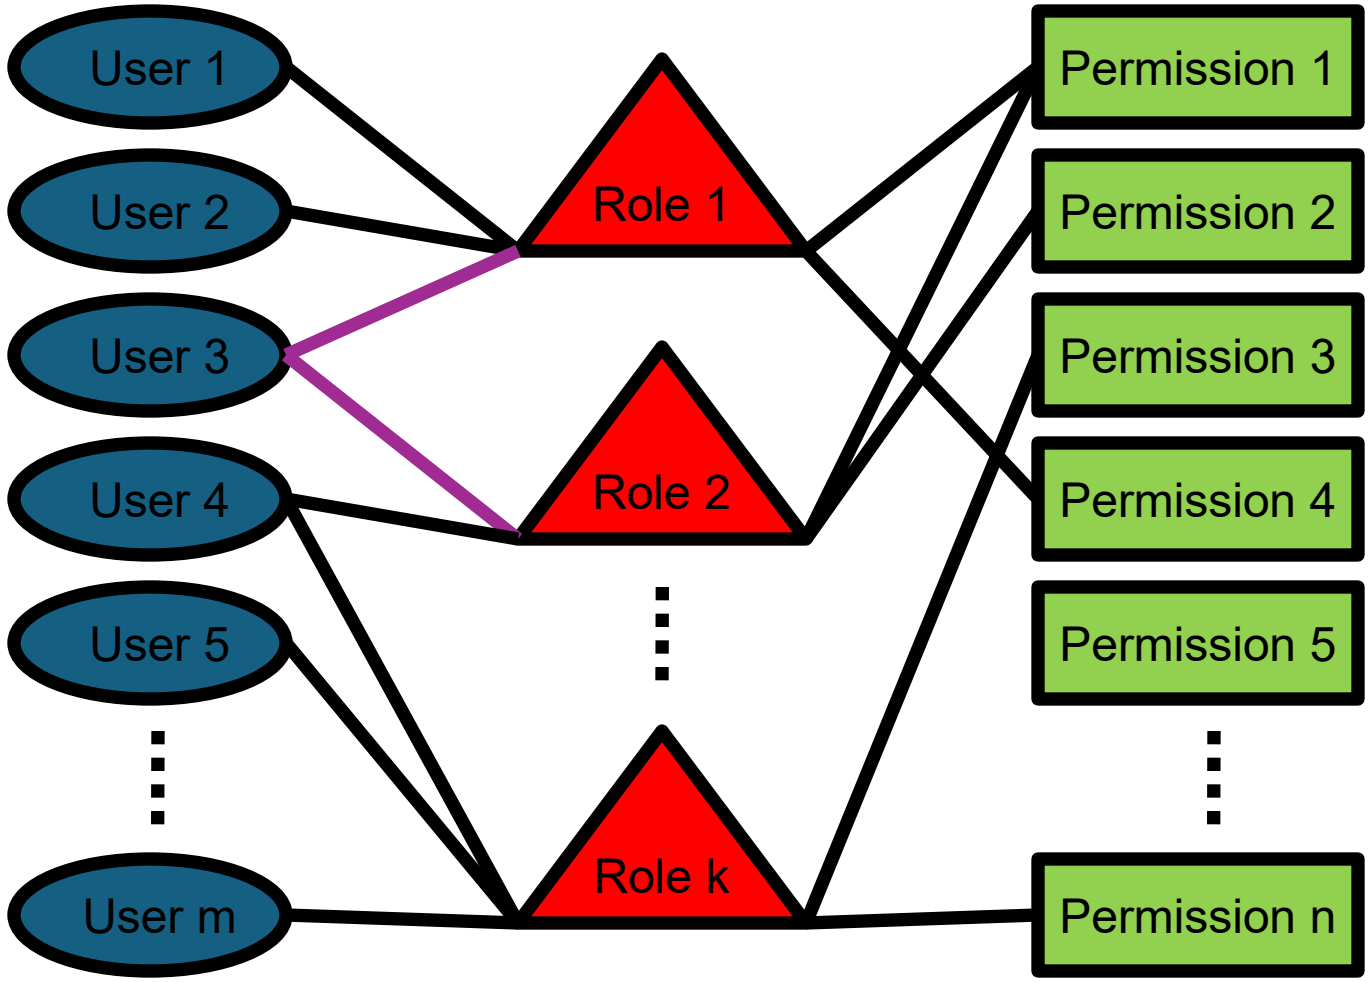

# A Formal RBAC System

- Defined over the following principals
  - U: user set
  - R: role set
  - P: permission set
  - S: session set (not always used)

- Relations
  - UA $\subseteq$ U $\times$ R (which users belong to which roles)
  - PA $\subseteq$ P $\times$ R (which permissions belong to which roles)
    - Note: Permissions are positive (not negative) statements

- Functions
  - User: S $\rightarrow$ U (e.g., session $s_i$ belongs to user $u_j$)
  - Roles: S $\rightarrow 2^{|R|}$ (mapping of each session to set of roles)

# “Core” RBAC Framework

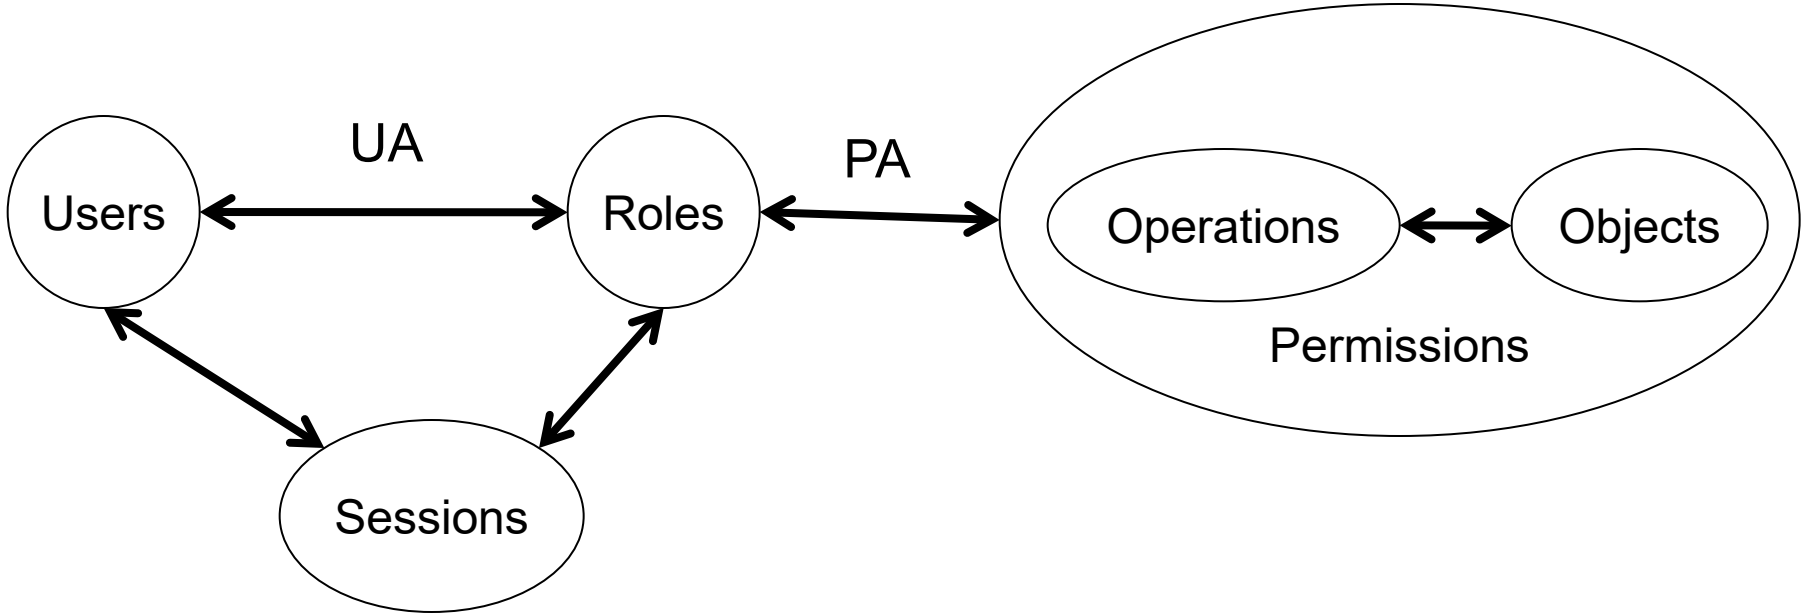


- Notice: permissions are often partitioned into Operations & Objects

D. Ferraiolo, et al. ACM Transactions on Information and System Security. 2001; 4(3): 224-274.

# Does RBAC Help?

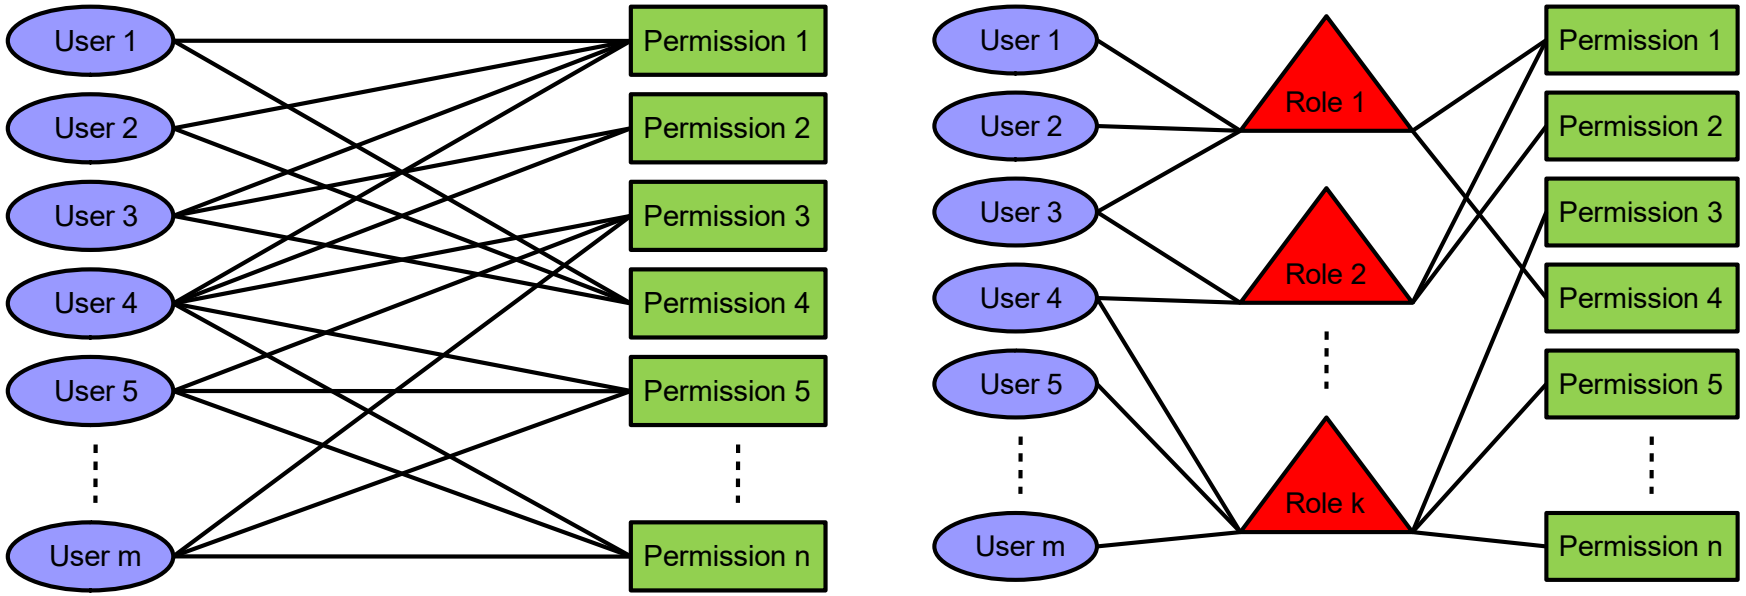

# RBAC in Practice

- Various database management systems (DBMS)*

- Enterprise Security Management
  - Take a look at the IBM Security Identity Governance and Intelligence (IGI)
  - https://www.ibm.com/us-en/marketplace/identity-governance-and-intelligence

- Various operating systems use RBAC in a limited way (think groups and rights)

# Readings for the Next Week

- 1. Kantarcioglu M, Jiang W, Liu Y, Malin B. **A cryptographic approach to securely share and query genomic sequences**. *IEEE Transactions on information technology in biomedicine*. 2008 Sep 3;12(5):606-17. https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=4358920

- Optional
  - ❑ 2. Jha S, Kruger L, Shmatikov V. **Towards practical privacy for genomic computation**. In *2008 IEEE Symposium on Security and Privacy (sp 2008)* 2008 May 18 (pp. 216-230). IEEE. https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=4531155
  - ❑ 3. 《Responsible Genomic Data Sharing Challenges and Approaches》 Chs.5&6.

# Feedback Survey

- One thing you learned or felt was valuable from today's class & reading

- Muddiest point: what, if anything, feels unclear, confusing or "muddy"

- https://www.wjx.cn/vm/hX0mIro.aspx