CSCI 1515 Applied Cryptography

This Lecture:

- Blockchain (Continued)
- Differential Privacy
- Privacy in ML

Blockchain

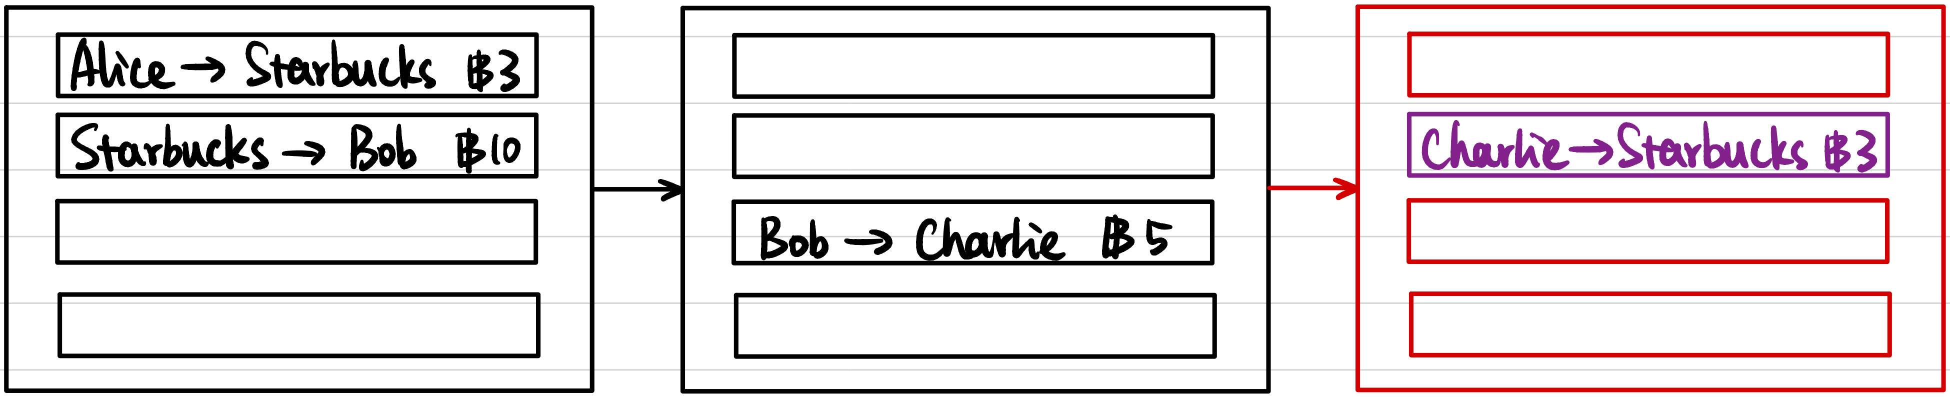

- Public ledger that everyone can view & verify
- Maintained by "miners" in a distributed way

Step 1: Charlie wants to make a transaction Charlie → Starbucks ฿3
↳ broadcasts it to the entire network

Step 2: All miners collect all transactions in the network
- Verify validity
  - ① initiated by sender ← How?
  - ② enough balance in sender's account
- Agree on next block
  - How?

Step 3: Repeat

Byzantine Agreement

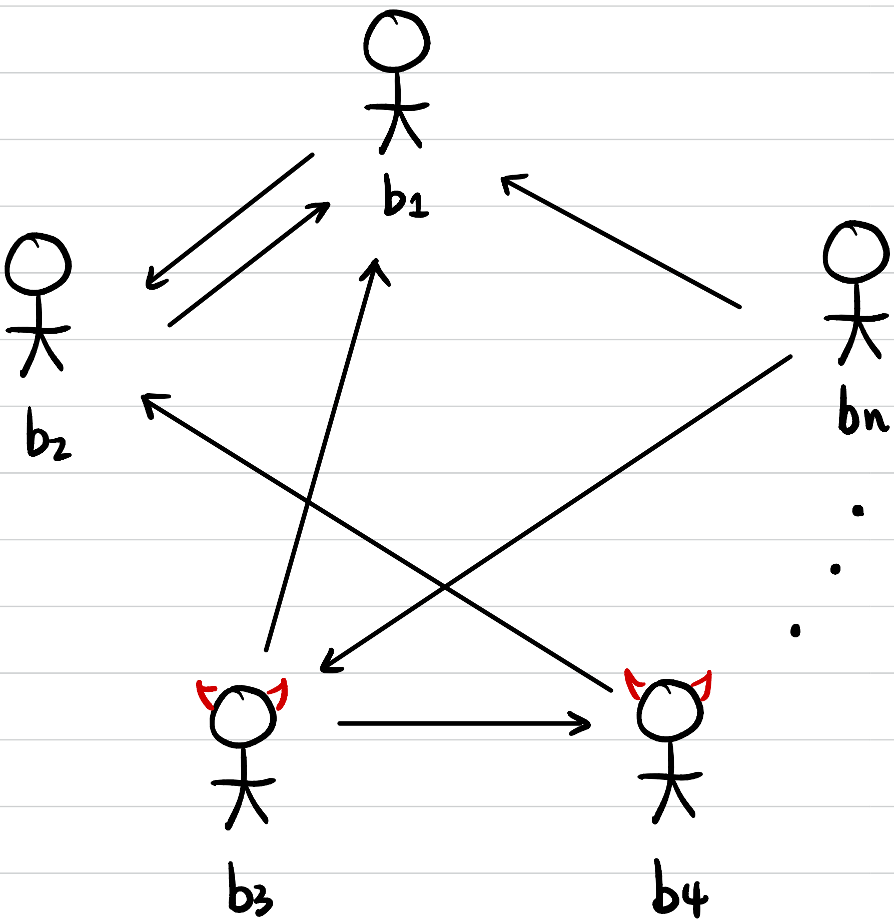

$\Downarrow$

Agree on a block

(Guaranteed Output Delivery)

Byzantine Fault Tolerance (BFT) Protocol:

If $n \geq 3t+1$, (necessary)

then it's possible to reach consensus.

Assume $t < n/3$, then agree on a valid block.

Any problem?

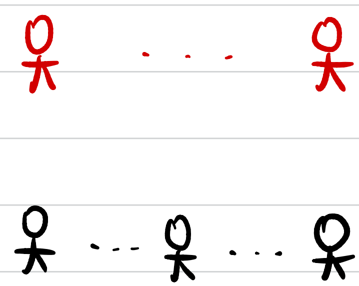

Sybil Attack

Proof of Work (PoW)

Miner 1:

Hash (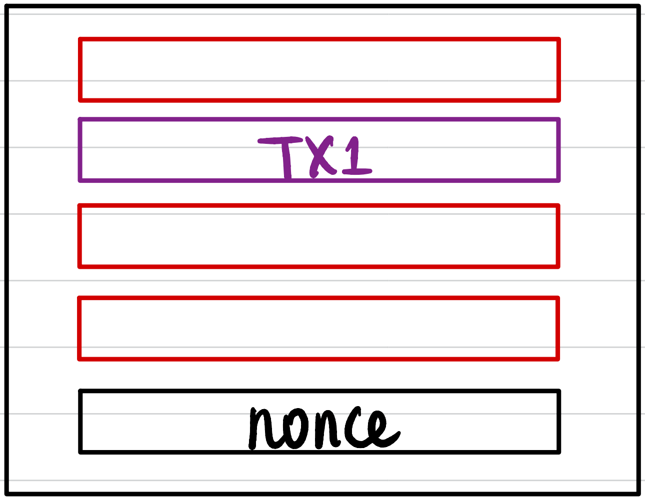

) $= \underbrace{00\cdots0}_{30}1011\cdots0$

Find nonce s.t. Hash(block) has ≥ 30 leading 0's.

Consensus Protocol:

Whoever first finds a block that hashes to a value w/ ≥ 30 leading 0's, that block becomes the next block.

Proof of Work (PoW)

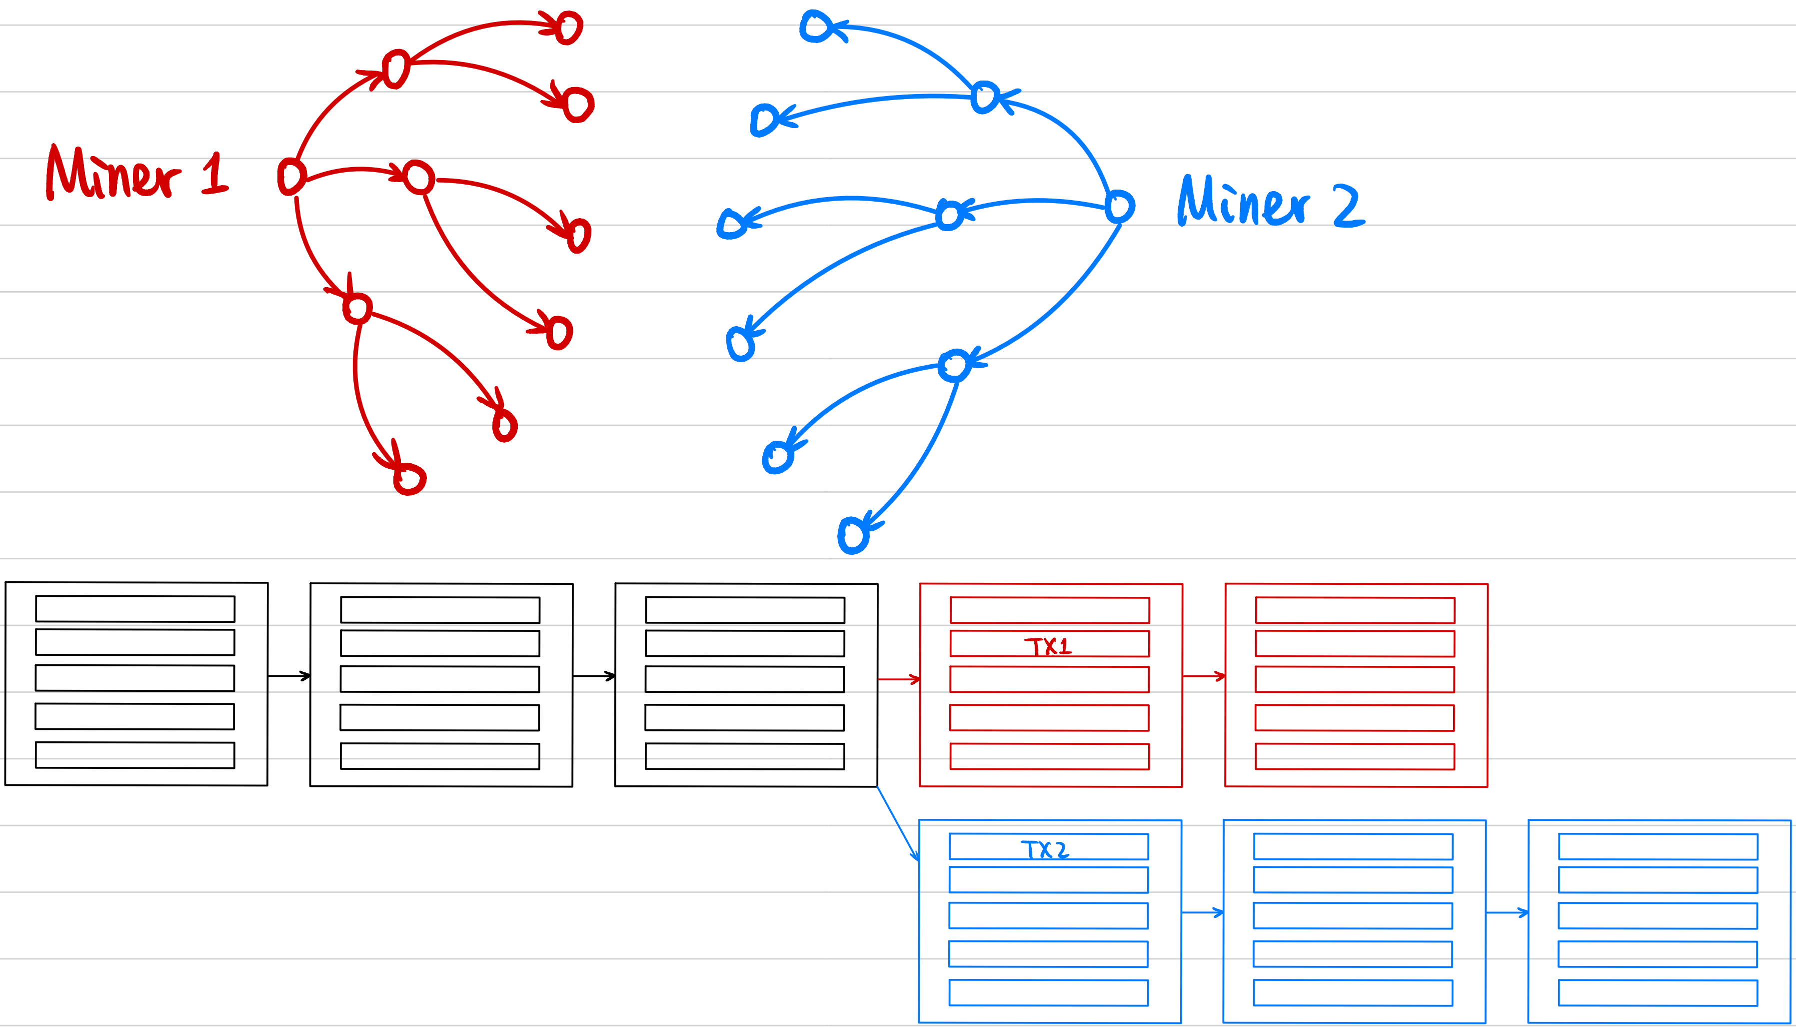

Longest Chain Rule: Always adopt the longest chain.

Assuming honest majority of computation power, the longest chain is always valid.

Extensions

- Smart contracts
- Proof of Stake (PoS)
- Anonymous transactions (zk-SNARGs)
- Public bulletin board

Differential Privacy

Name	Age	Gender	Race	Weight	ZIP	Disease
Alice						
Bob						
Charlie						
David						
Emily						
Fiona						

Want to make the (sensitive) data public / available to others (e.g. for medical study).

Attempt 1: "Anonymize" the data.
Delete personally identifiable information (PII): name, DOB, ...

Attempt 2: Only answer aggregate statistics queries.

Privacy Guarantee?

Access to the output shouldn't enable one to learn ~~anything~~ much more about an individual compared to one ~~without access.~~ with access to the output computed on a database without the individual.

Is this possible?

Privacy vs. Utility

Differential Privacy

x_1
x_2
$\vdots$
x_n

$$D \in X^n \longrightarrow \boxed{M} \longrightarrow M(D)$$

Def ε-Differential Privacy for a randomized mechanism:

$\forall$ neighboring datasets D_1 & D_2 (differing in one row),

$\forall\, T \subseteq \text{range}(M)$,

$$\Pr[M(D_1) \in T] \leq e^{\varepsilon} \cdot \Pr[M(D_2) \in T]$$

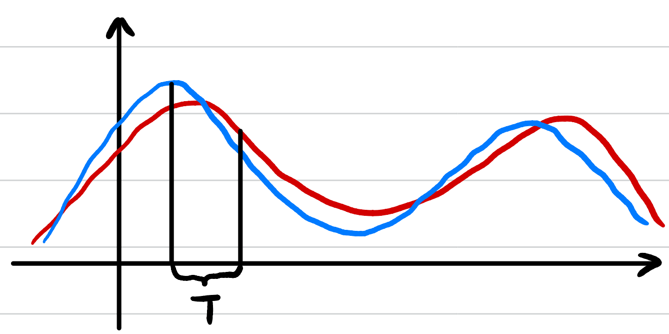

Differential Privacy

Def (ε, δ) - Differential Privacy for a randomized mechanism:

$\forall$ neighboring datasets D_1 & D_2 (differing in one row),

$\forall T \subseteq \text{range}(M)$,

$$\Pr[M(D_1) \in T] \leq e^{\varepsilon} \cdot \Pr[M(D_2) \in T] + \delta$$

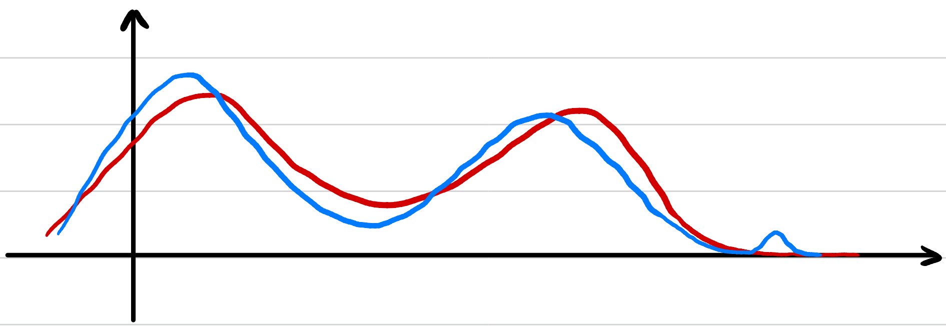

Is a bigger ε better for privacy, or worse? Worse

Is a bigger δ better for privacy, or worse? Worse

Randomized Response

Counting query: What percentage $\nearrow \alpha$ of individuals satisfy predicate P?

For each row x_i:

① Sample $b \xleftarrow{\$} \{0,1\}$

② If $b=0$, then $y_i := P(x_i)$

Otherwise, $y_i \xleftarrow{\$} \{0,1\}$

$M(D) := (y_1, y_2, \cdots, y_n)$

Thm Randomized Response is $\ln 3$-DP.

How to make the mechanism more private? Flip a biased coin in ①

How to estimate the query output?

$$\mathbb{E}[\# 1\text{'s}] = \frac{1}{2}\cdot\alpha\cdot N + \frac{1}{2}\cdot\frac{1}{2}\cdot N \approx \frac{k}{N}$$

Laplace Mechanism

Def Sensitivity of a function $f: X^n \to \mathbb{R}$

$$\Delta f := \max_{D_1 \sim D_2} |f(D_1) - f(D_2)|$$

Laplace Mechanism: $M(D) = f(D) + \mathrm{Lap}(\Delta f / \varepsilon)$

Thm The Laplace Mechanism is ε-DP.

Laplace distribution:

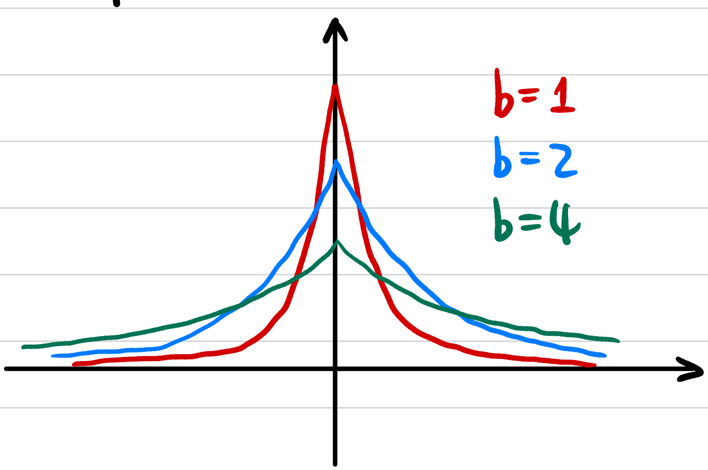

$\mathrm{Lap}(b)$:

$$\mathrm{PDF}(x) = \frac{1}{2b} \cdot \exp\left(-\frac{|x|}{b}\right)$$

For $X \sim \mathrm{Lap}(b)$, $\Pr[|X| \geq bt] \leq \exp(-t)$

Is a bigger b better for privacy, or worse?

Better

Composition Theorems

Thm (post-processing) If $M: X^n \to Y$ is (ε, δ)-DP,

$f: Y \to Z$ is an arbitrary randomized function,

then $f \circ M: X^n \to Z$ is also (ε, δ)-DP.

Thm (group privacy) If $M: X^n \to Y$ is $(\varepsilon, 0)$-DP,

then M is $(k \cdot \varepsilon, 0)$-DP for groups of size k.

Thm (composition) If $M_i: X^n \to Y$ is $(\varepsilon_i, \delta_i)$-DP $\forall i \in [k]$,

then $M(D) := (M_1(D), \cdots, M_k(D))$ is $\left(\sum_{i \in [k]} \varepsilon_i, \sum_{i \in [k]} \delta_i\right)$-DP.

Privacy in ML

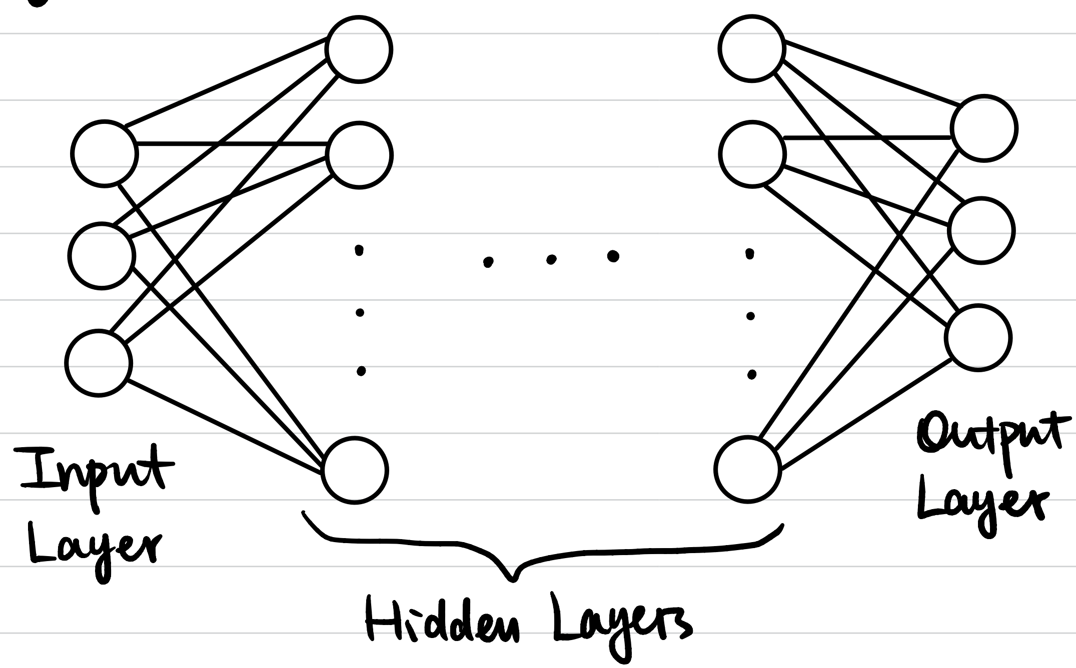

Each node in hidden layers: linear function + activation function

Data points $(\vec{x}_i, y_i)$

ML model: weights $\vec{w}$

Loss function $L_i(\vec{w})$

Stochastic Gradient Descent (SGD):

- $\vec{w}$ initialized randomly
- Each iteration:

$$\vec{w} \leftarrow \vec{w} - \eta \cdot \nabla L_i(\vec{w})$$

$$\vec{w} \leftarrow \vec{w} - \frac{\eta}{B} \cdot \sum_{i \in [B]} \nabla L_i(\vec{w})$$

Privacy in ML

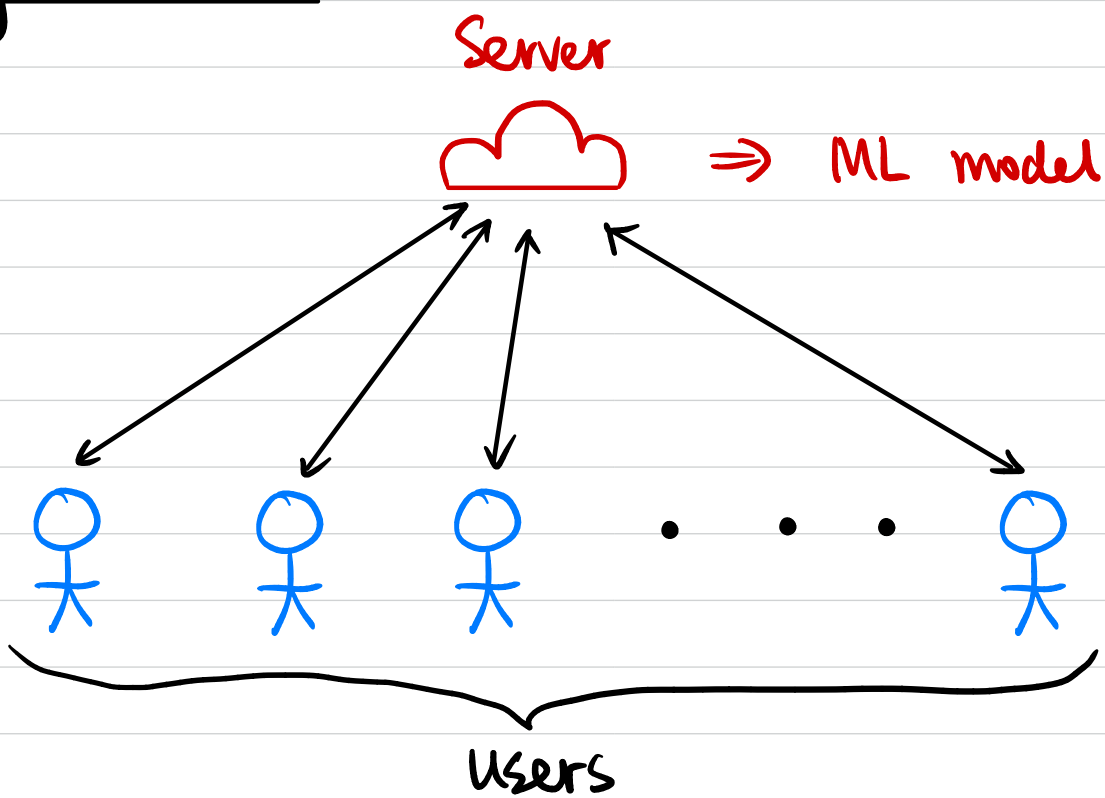

- Does the model (updates) contain private information?
- Secure inference / training?
- Data deletion from trained model?