

Dordis

Efficient Federated Learning

with Dropout-Resilient Differential Privacy

Zhifeng Jiang,
Wei Wang,
Ruichuan Chen

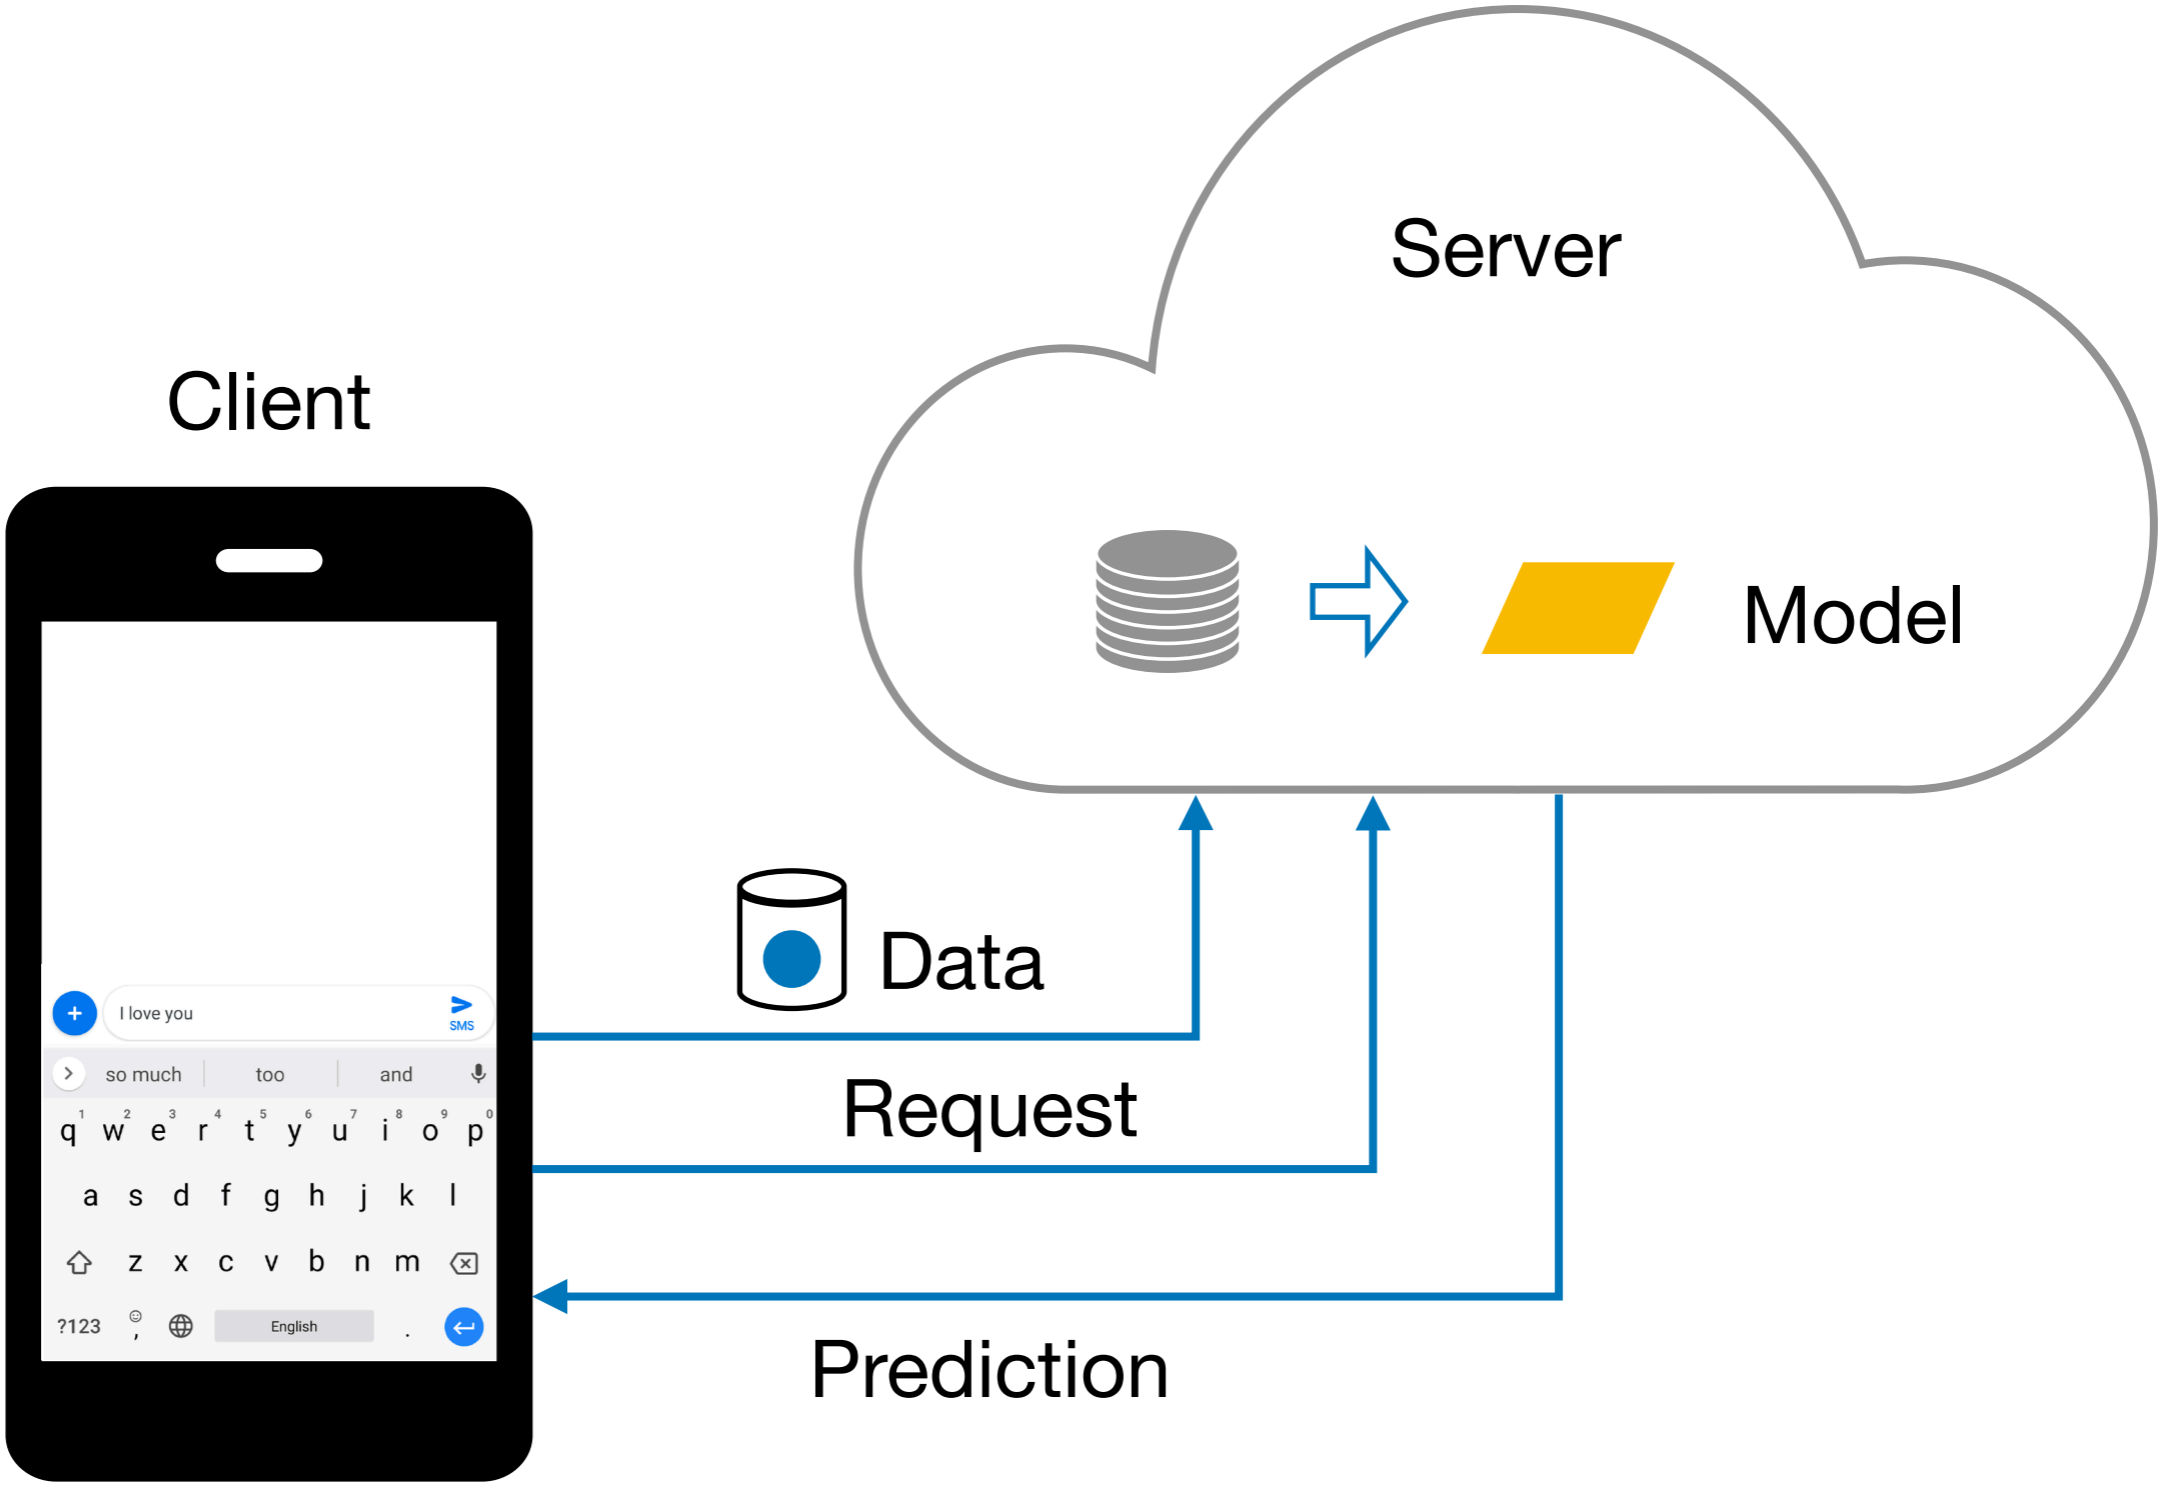


香港科技大學

THE HONG KONG UNIVERSITY OF
SCIENCE AND TECHNOLOGY

NOKIA
BELL
LABS

Centralized learning




Centralized learning **hurts privacy**

Forbes

Data breaches...

**Clearview AI, The Company
Whose Database Has Amassed 3
Billion Photos, Hacked**



Centralized learning **hurts privacy**

Data breaches...

Forbes

Clearview AI, The Company
Whose Database Has Amassed 3
Billion Photos, Hacked

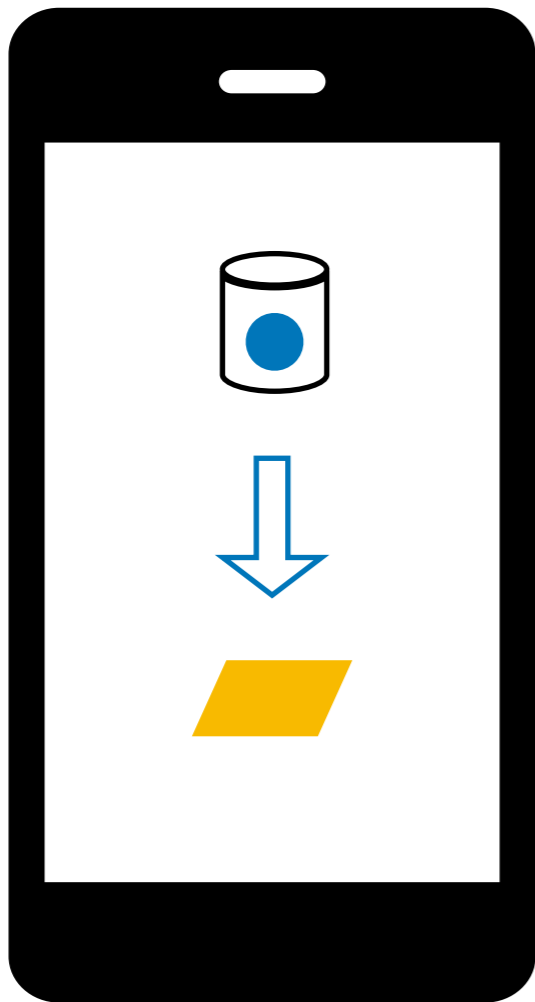
Potential abuse...

theguardian


Facebook halts use of WhatsApp data
for advertising in Europe

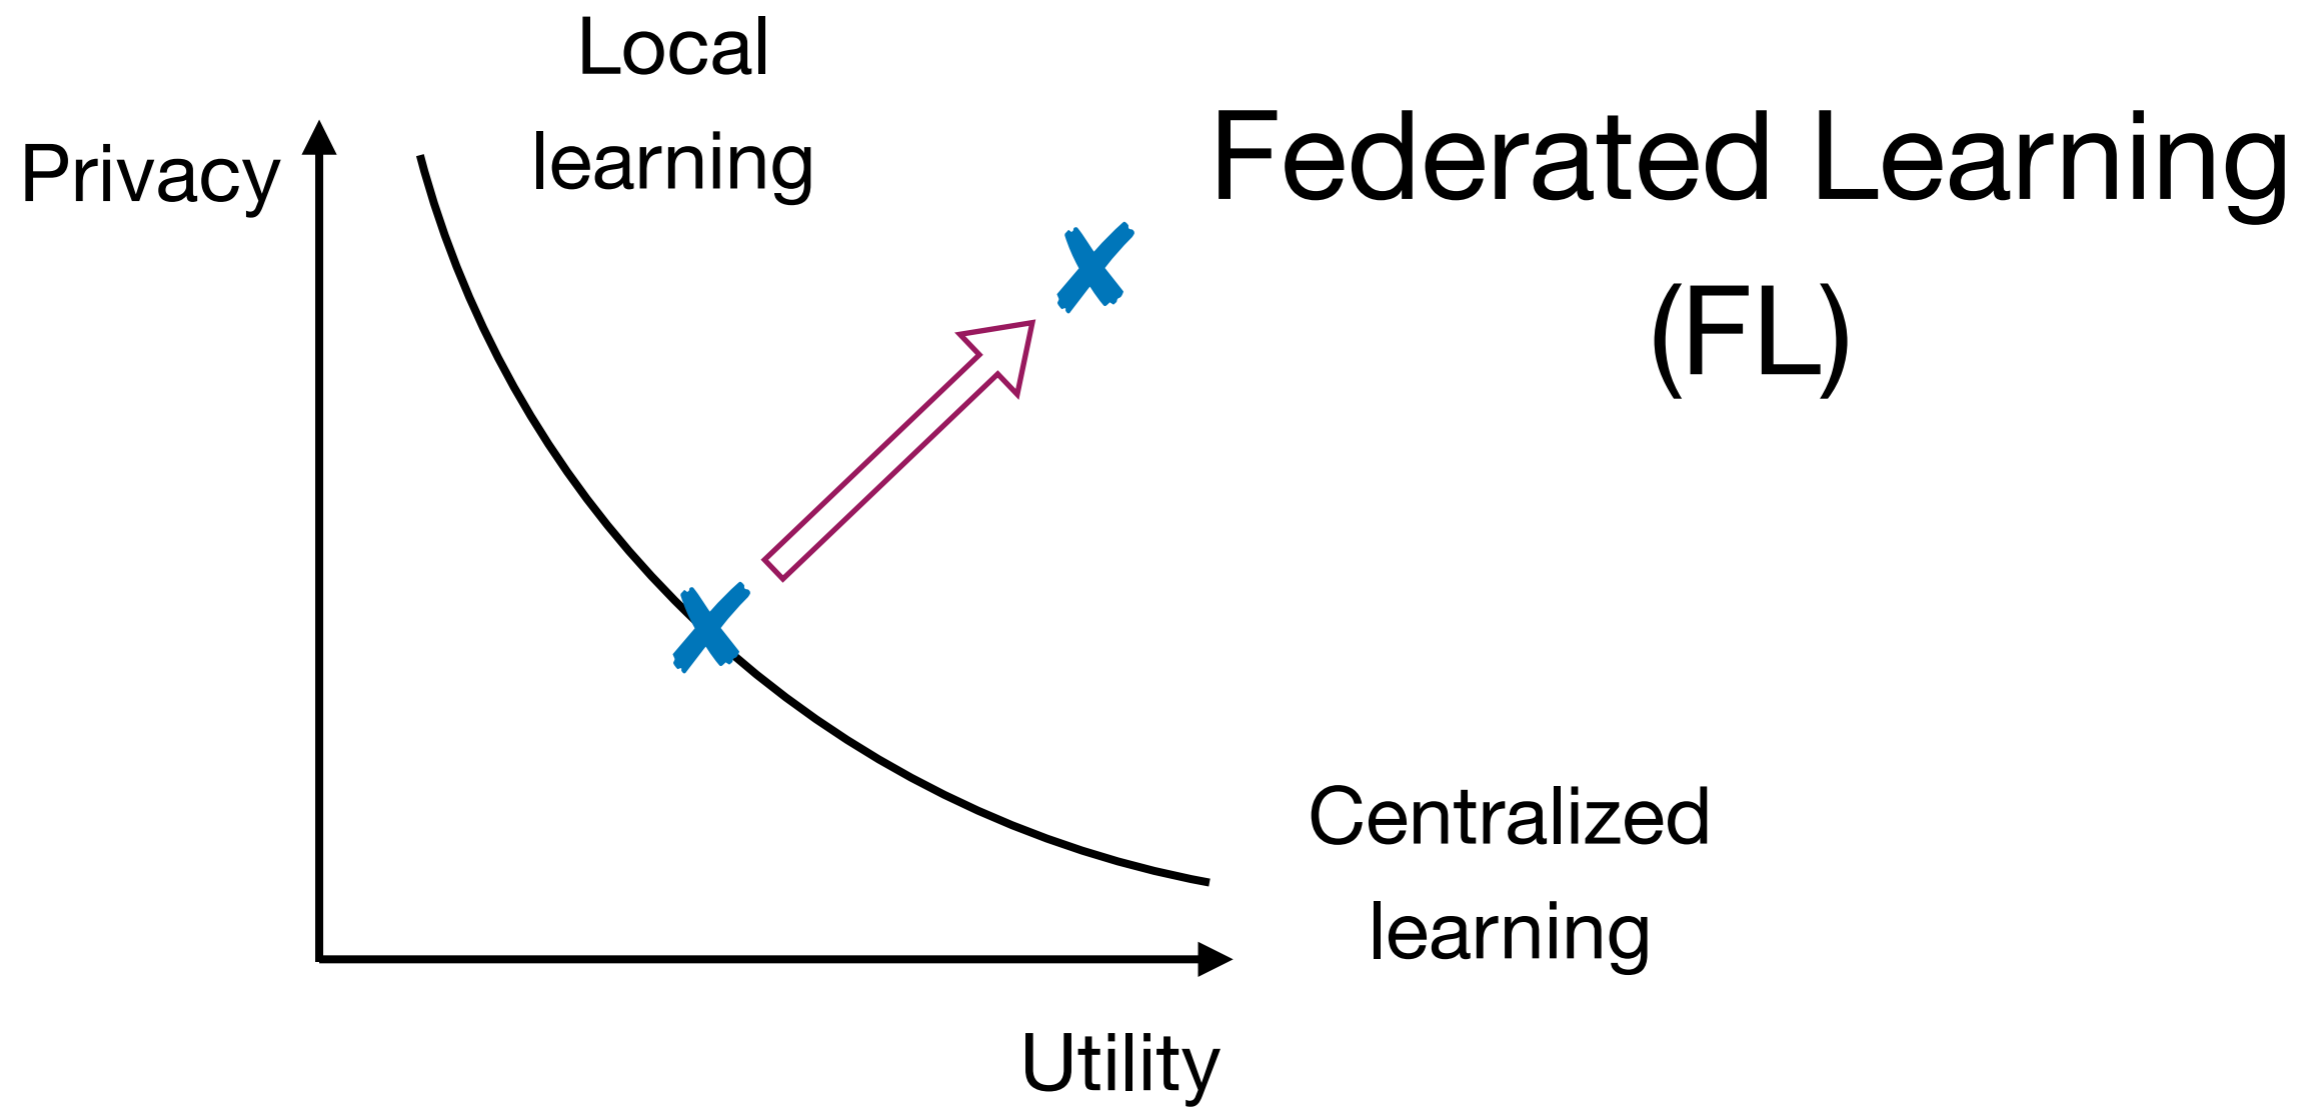


Local learning

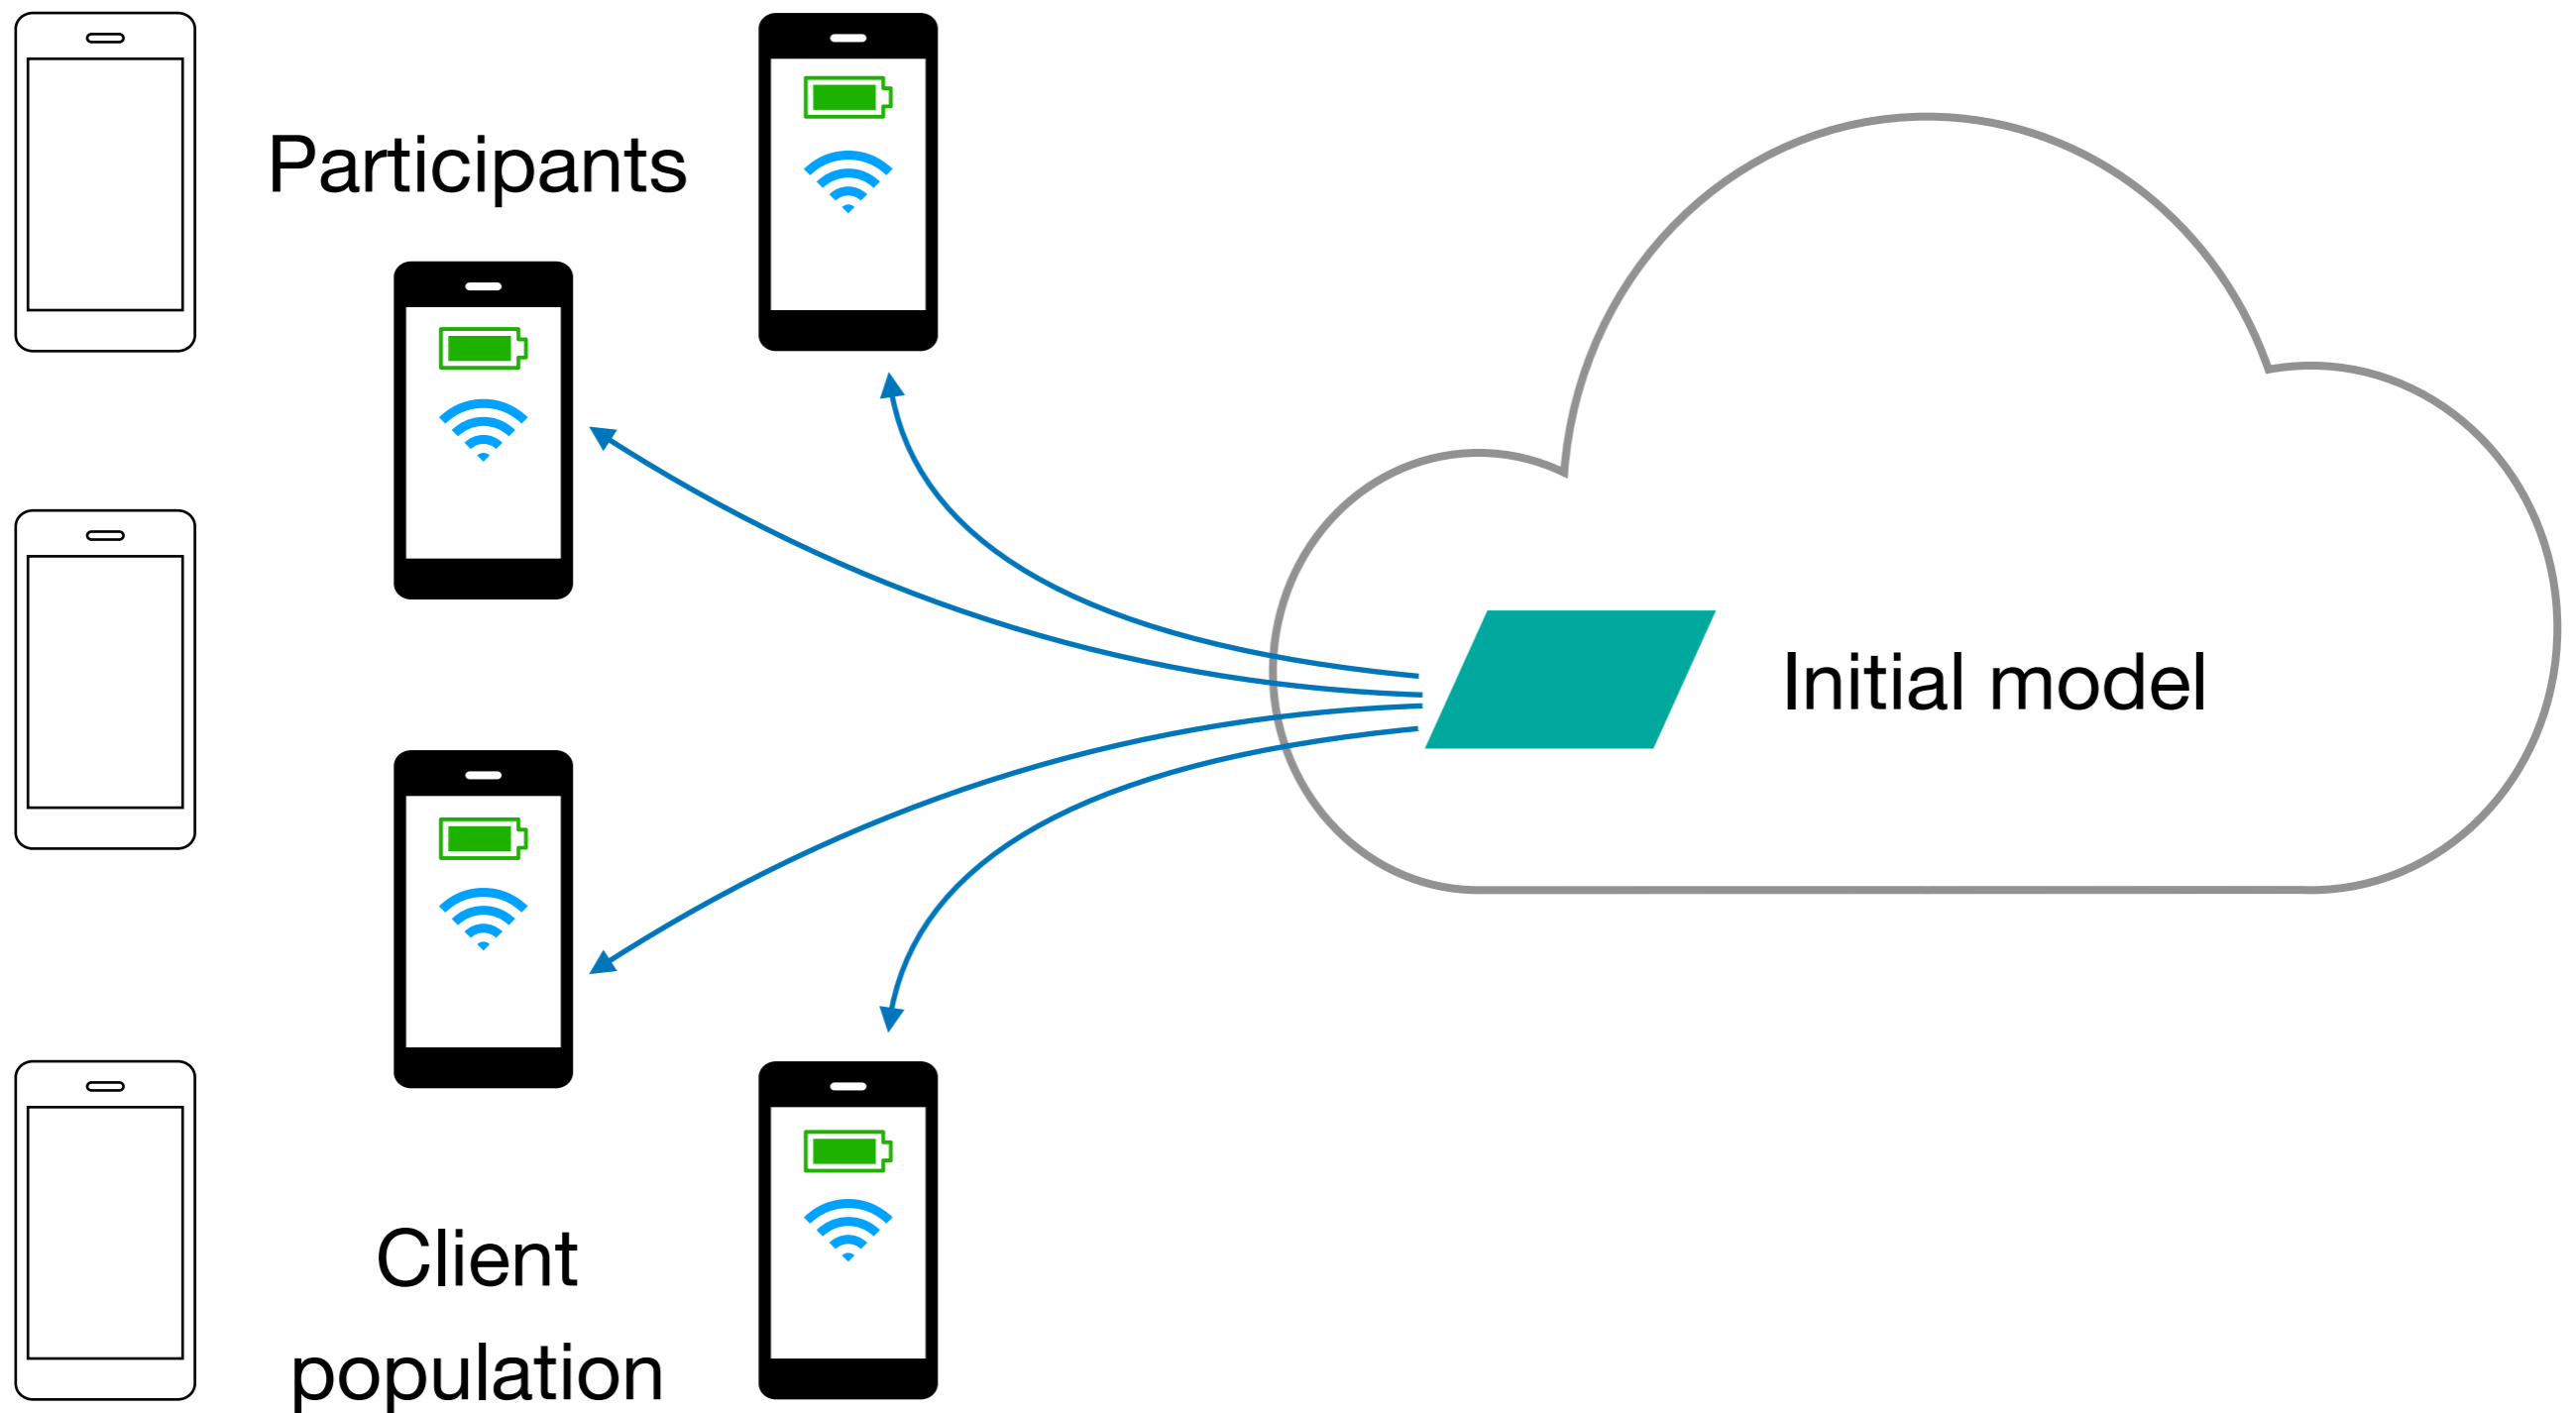


Local learning suffers from **low utility**






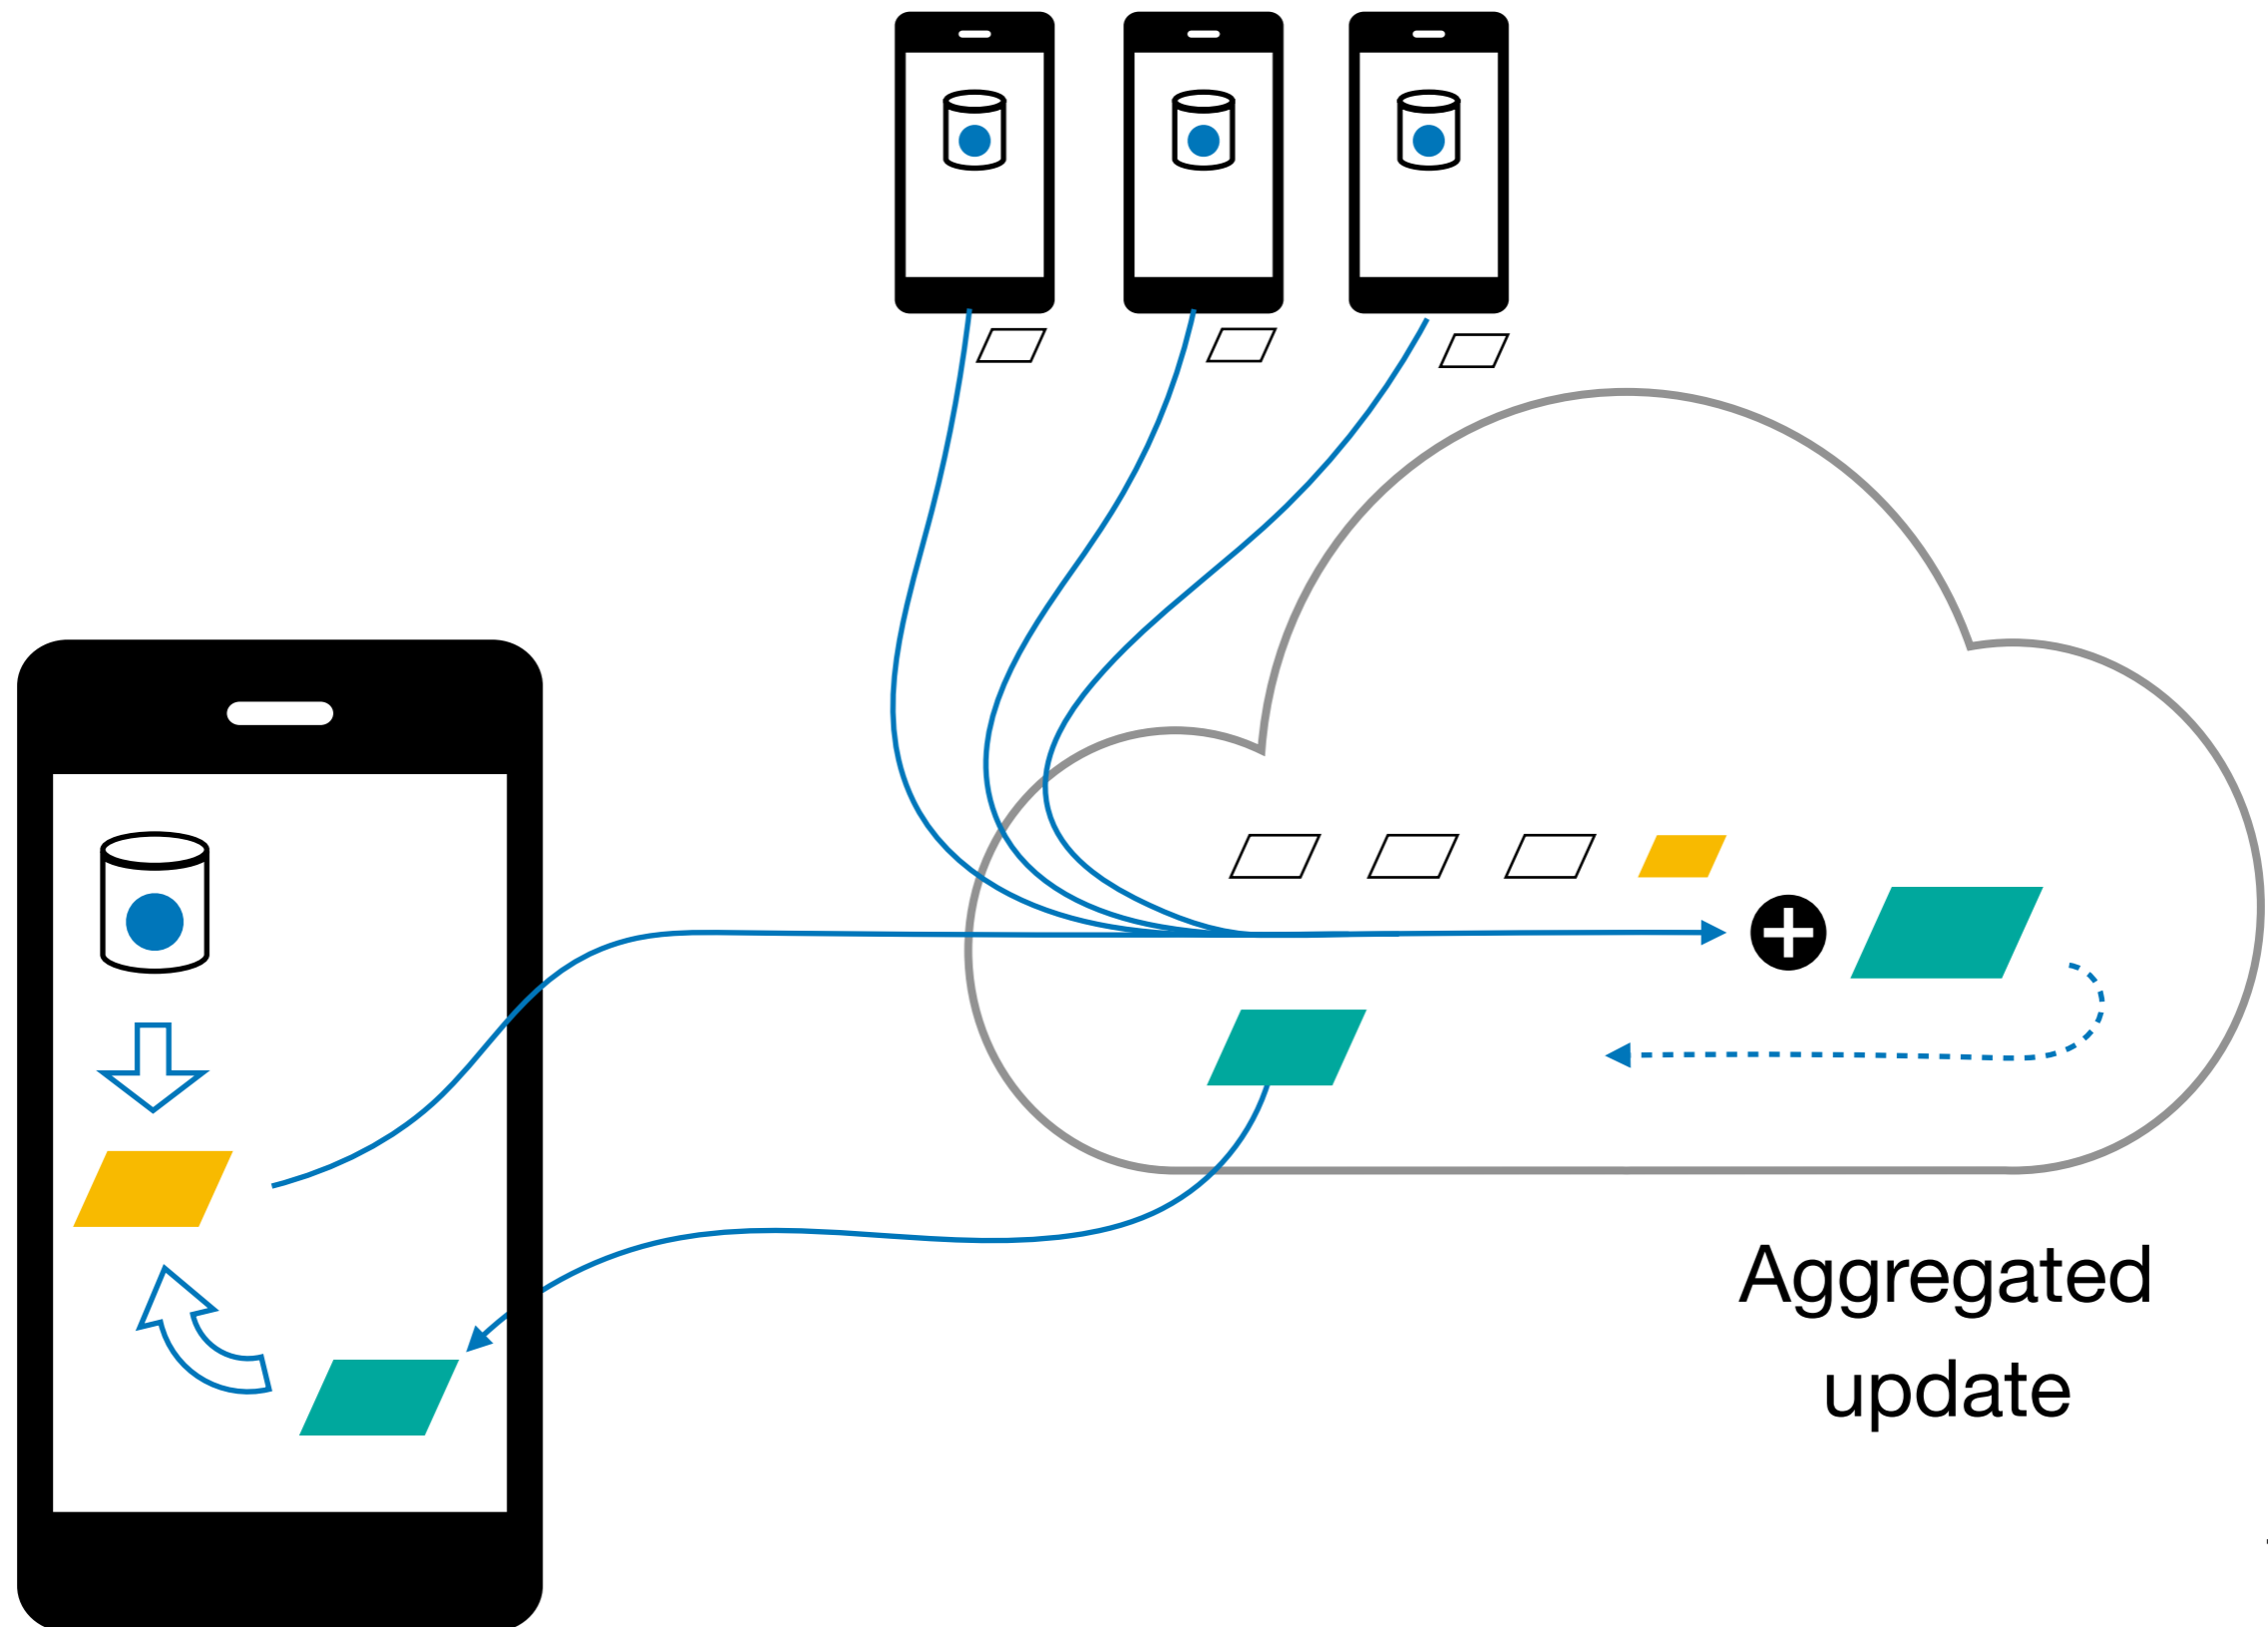
FL Step 1: Participant Selection



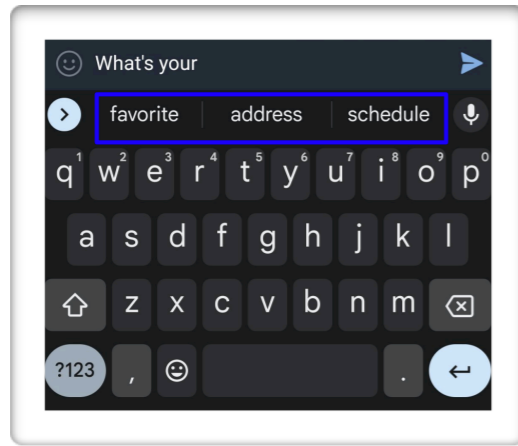
FL Step 2: Local Training



FL Step 3: Model Aggregation



FL Applications

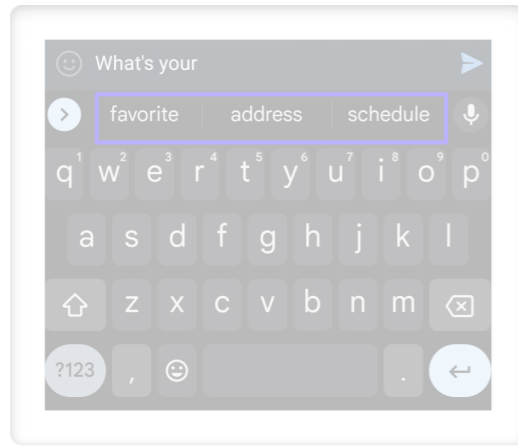


Mobile

Google's Keyboard

FL Applications

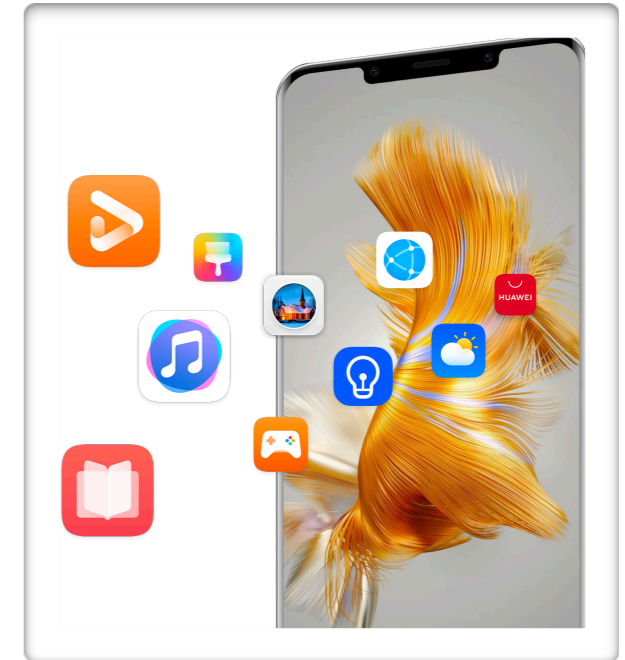
Mobile




Google's Keyboard



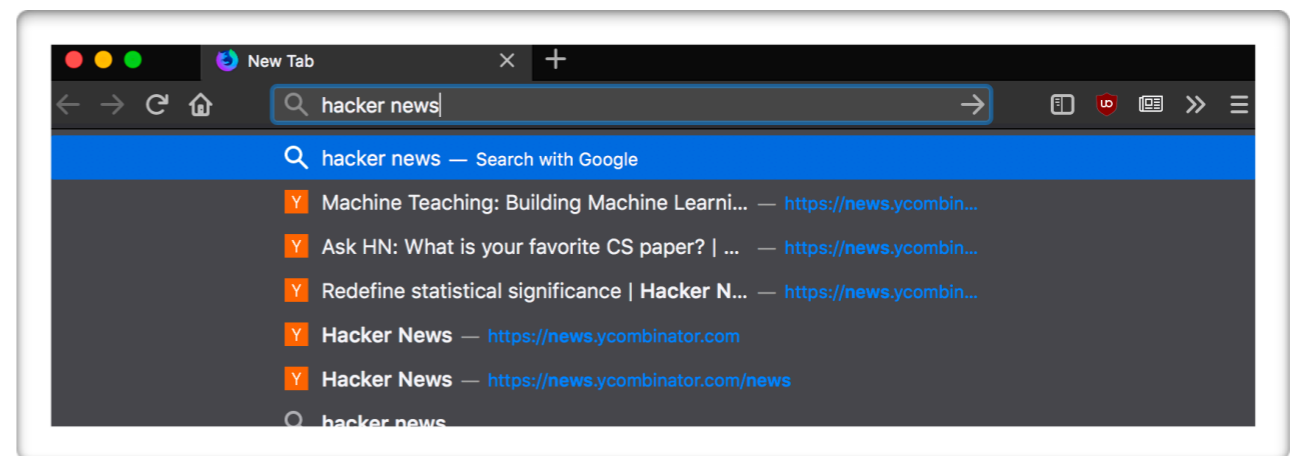
Apple's speaker recognition



Huawei's ads recommendation




Brave's news recommendation




Firefox's URL bar suggestion


IoT



Volvo's trajectory prediction

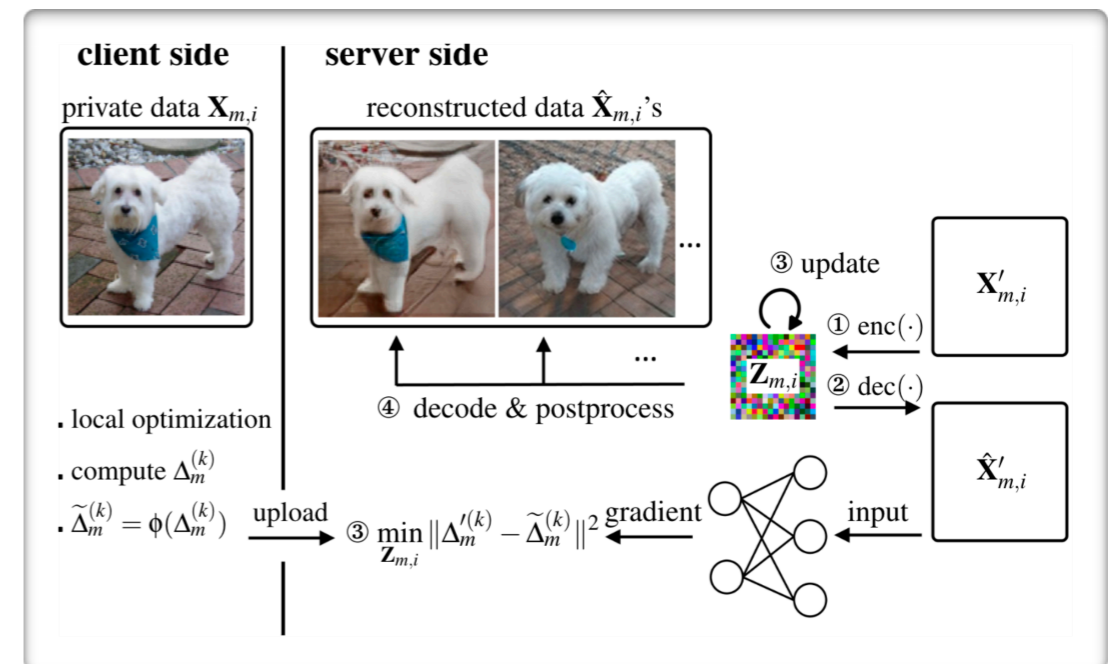



Cisco's 3D printing



Leveno's clogging detection

Data Leakage Remains in FL



e.g., data reconstruction¹ (Security '23)

[1] Gradient Obfuscation Gives a False Sense of Security in Federated Learning

Data Leakage Remains in FL

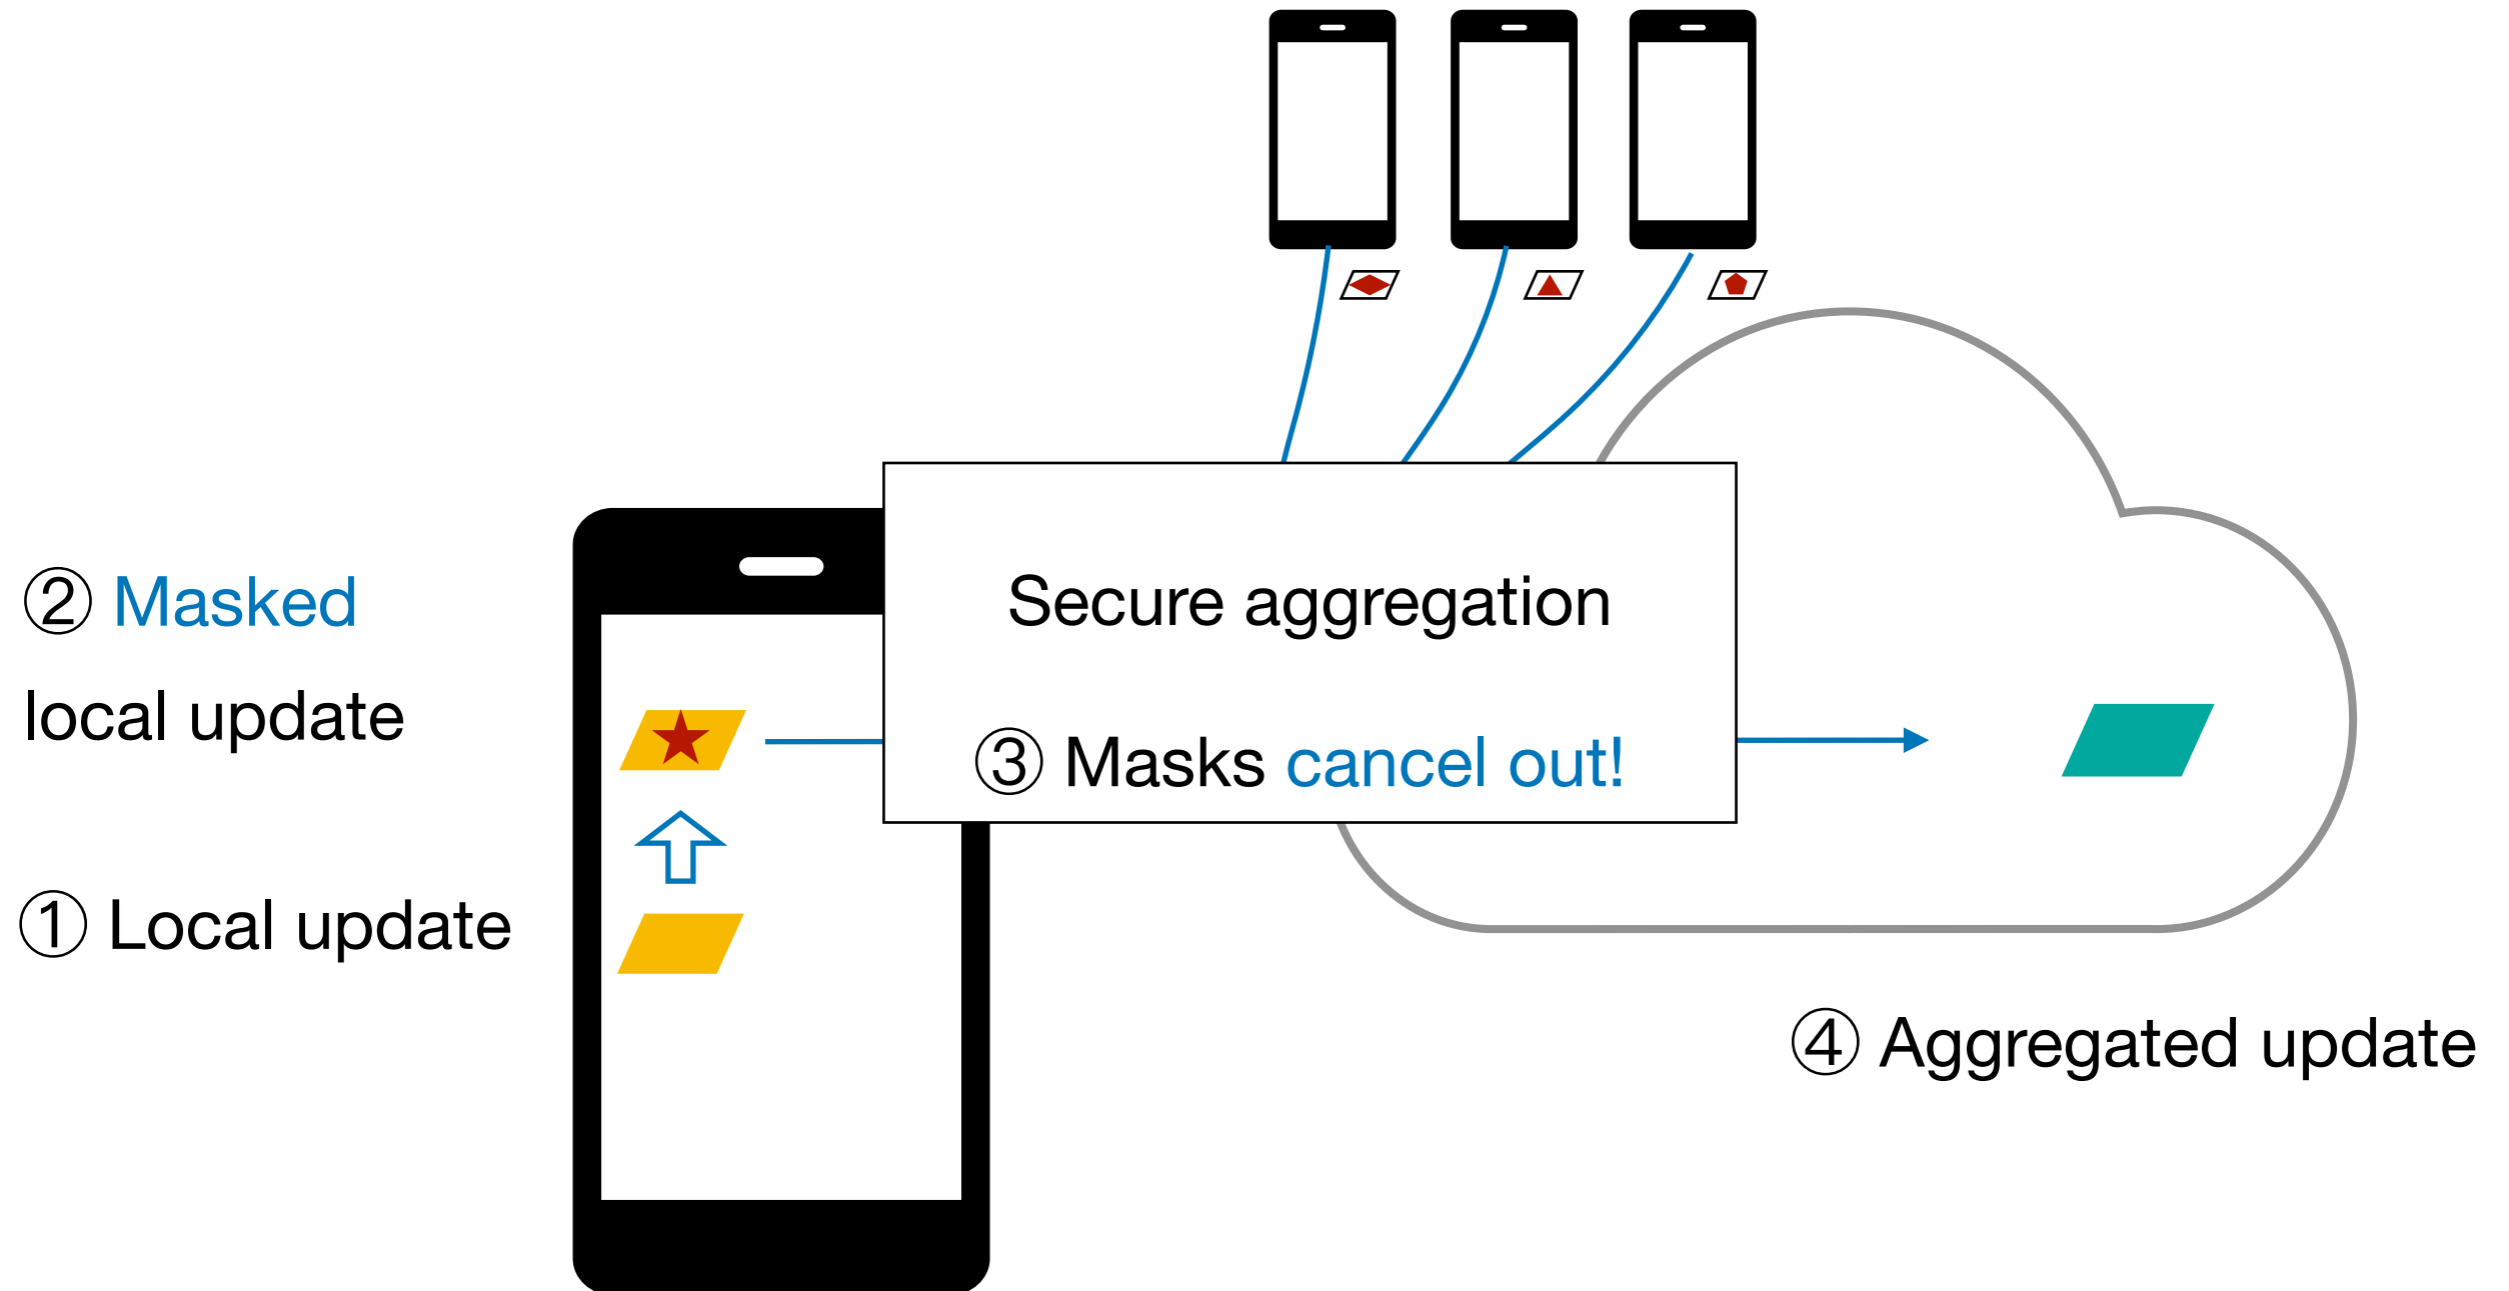
To conceal **local** updates?

Secure aggregation¹²
(CCS '17, '20)

[1] Practical secure aggregation for privacy-preserving machine learning
[2] Secure Single-Server Aggregation with (Poly) Logarithmic Overhead

Data Leakage Remains in FL

To conceal **local** updates?



Data Leakage Remains in FL

To conceal local updates?

To also perturb the **aggregated** update?

Sacrifice the
precision

Differential Privacy¹

For enhanced
privacy

Data Leakage Remains in FL

To conceal local updates?

To also perturb the **aggregated** update?

Sacrifice the
precision





Differential Privacy¹

For enhanced
privacy



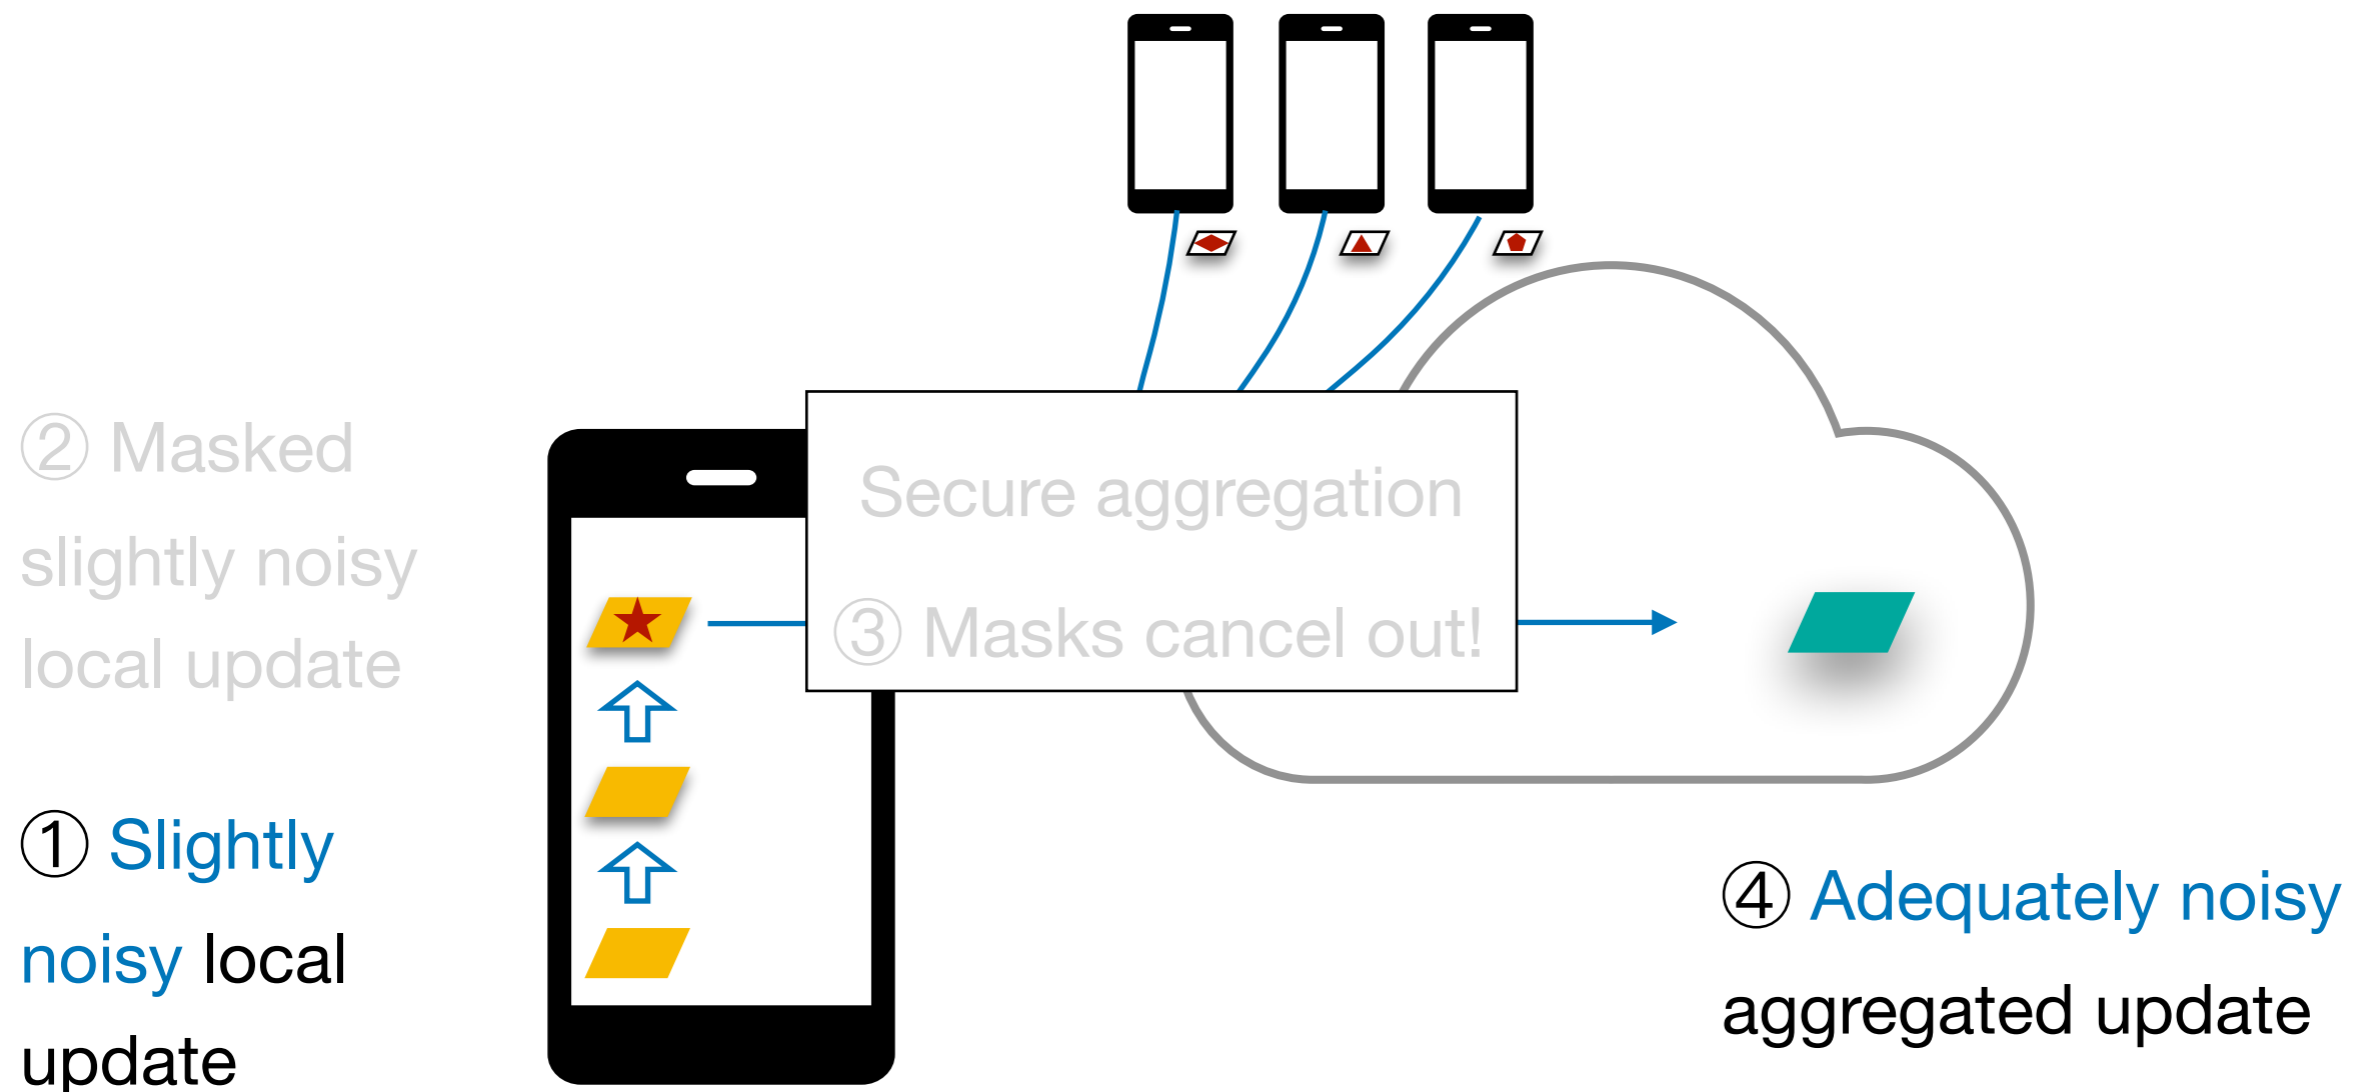
$$\begin{aligned} \text{aggregation} &= A(\text{local updates}) \\ &= f(\text{local updates}) + \text{random noise} \end{aligned}$$

DP ensures that  be **insensitive** to the impact of any single local update in 

Data Leakage Remains in FL

To conceal local updates?

To also perturb the **aggregated** update?

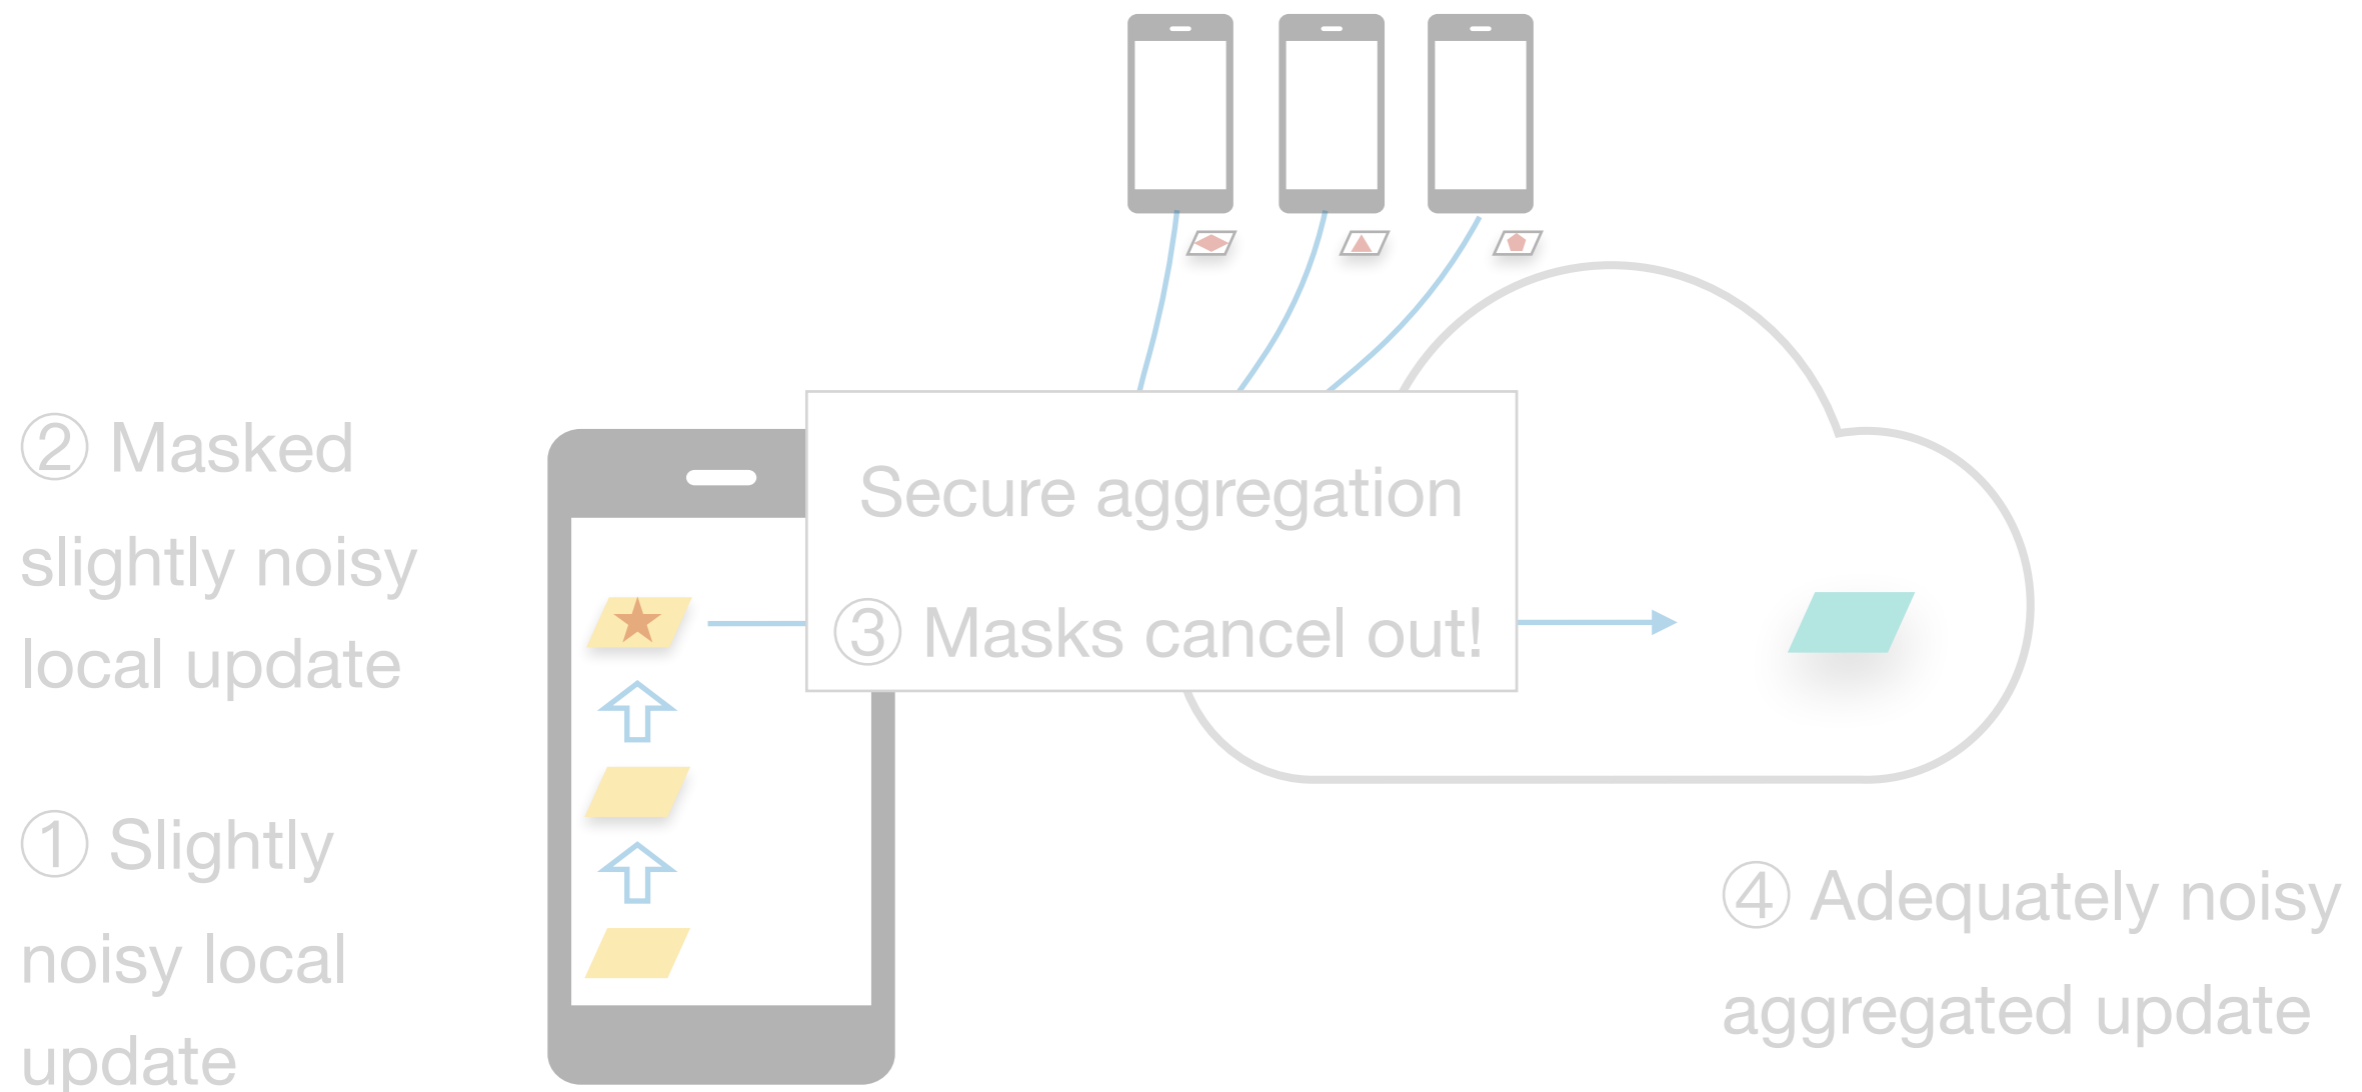


① Global privacy budget $\epsilon \rightarrow$ Calculate the **minimum required noise**

Distributed DP = SecAgg + DP

To conceal local updates?

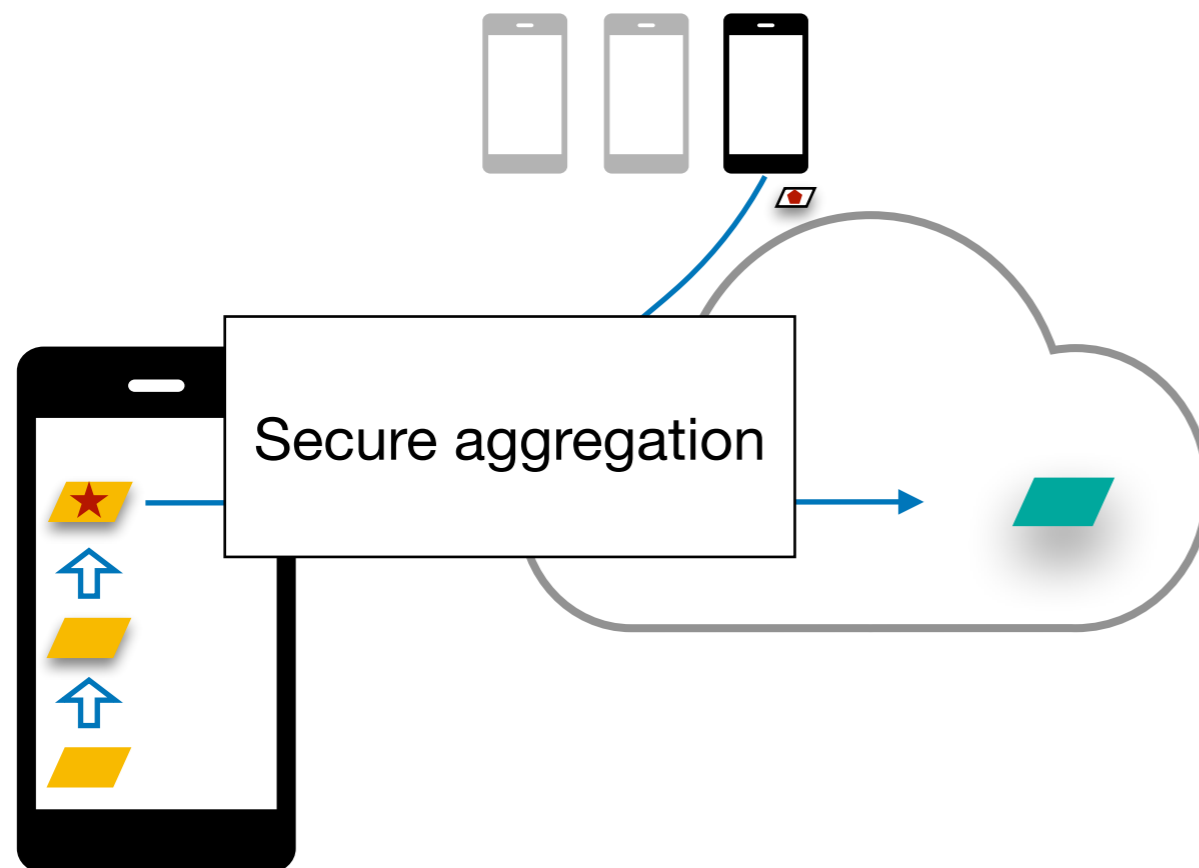
To also perturb the aggregated update?



① Global privacy budget $\epsilon \rightarrow$ Calculate the minimum required noise

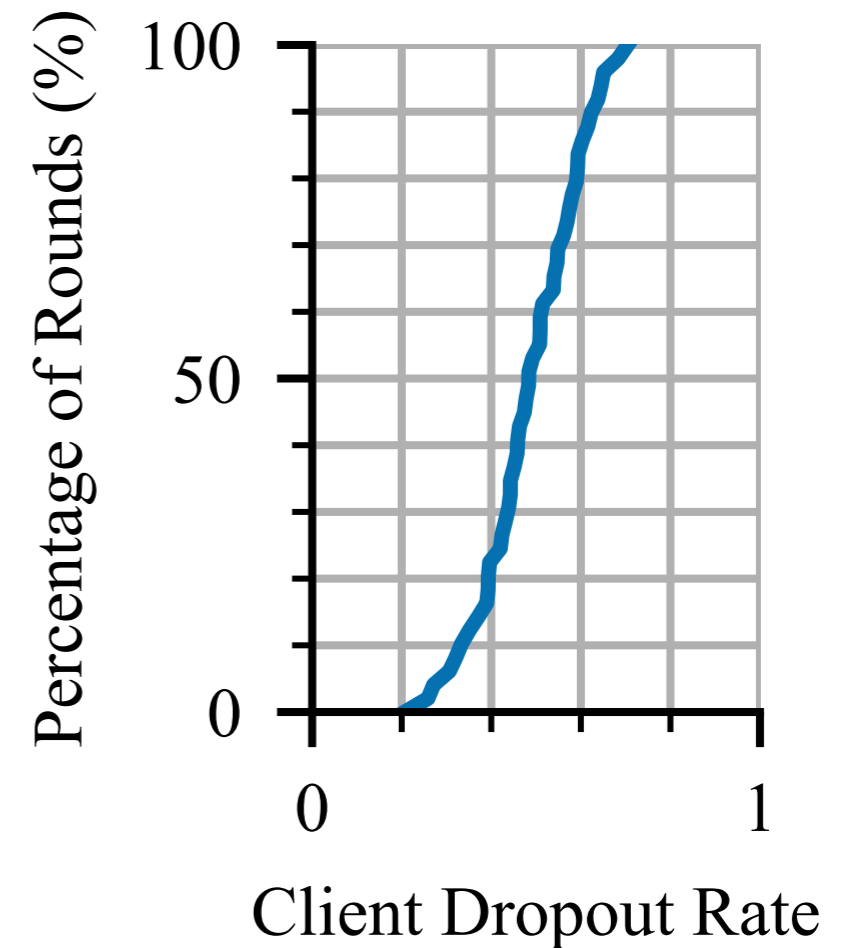

Distributed DP Has Two Practical **Issues**

1. **Privacy** Issue: caused by **client dropout**



Privacy Issue Caused by Client Dropout

Client dropout can occur anytime



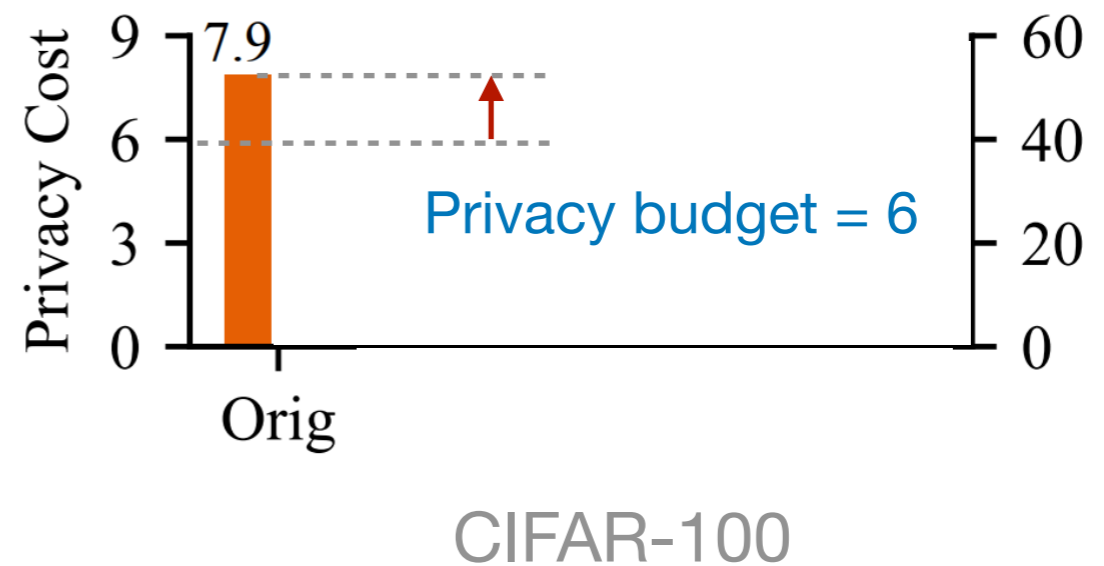
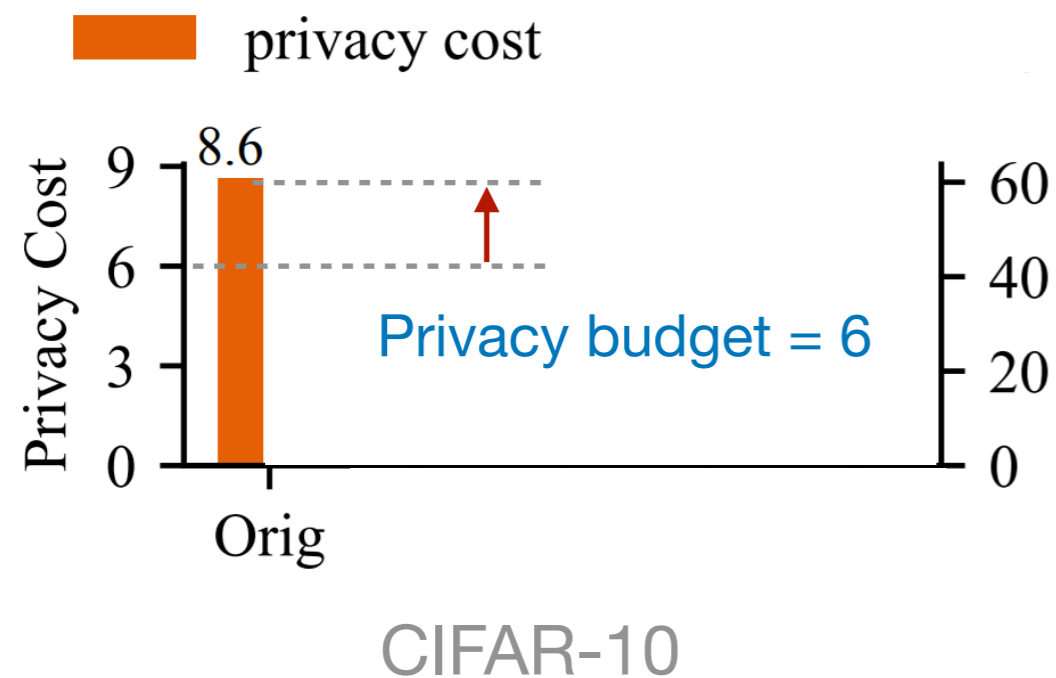
Client behaviors simulated with 100 volatile users from the FLASH dataset¹ (WWW '21)

[1] Characterizing impacts of heterogeneity in federated learning upon large-scale smartphone data

Privacy Issue Caused by Client Dropout

Client dropout can occur anytime

Insufficient noise for target privacy



Dropout-Resilient Noise Enforcement


Goal: *always* enforce the *target* noise level

Dropout-Resilient Noise Enforcement

Goal: always enforce the target noise level

Intuition: **add**-then-remove

- Each client first adds excessive noise as **separate components**

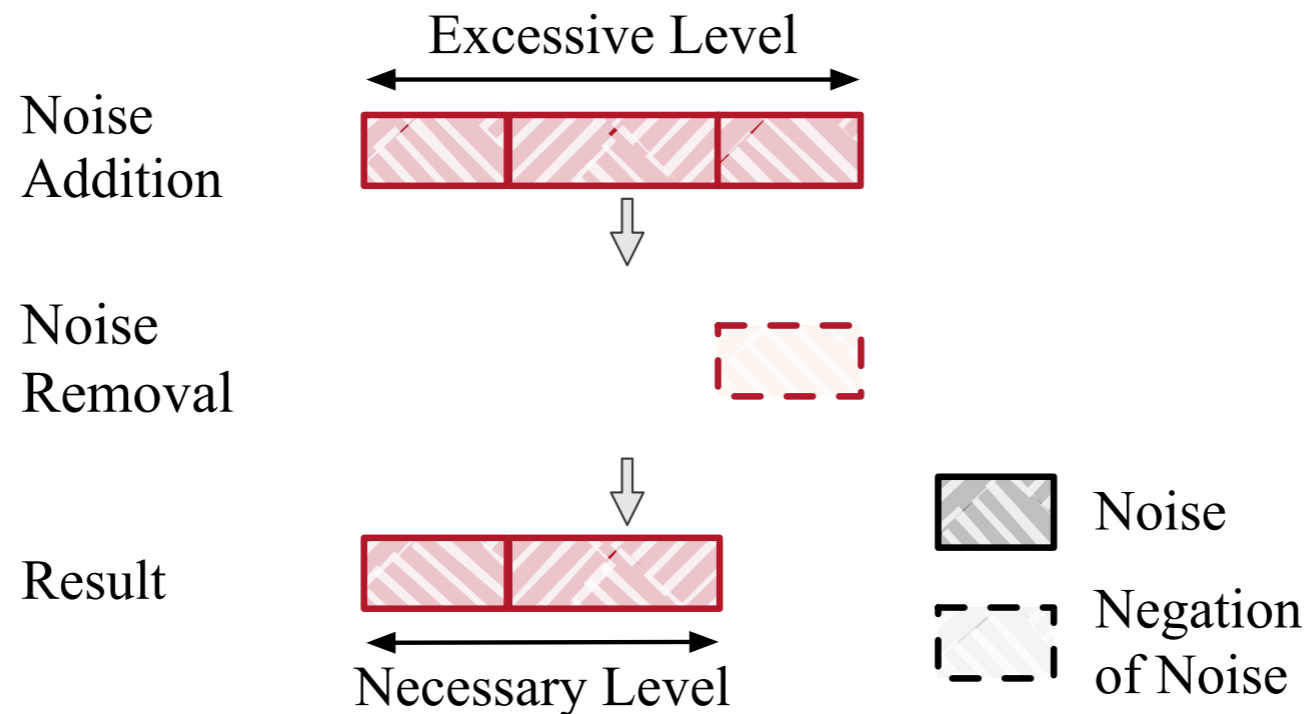


Dropout-Resilient Noise Enforcement

Goal: always enforce the target noise level

Intuition: add-then-remove

- Each client first adds excessive noise as separate components
- After aggregation, unnecessary ones are removed by the server



Dropout-Resilient Noise Enforcement

Goal: always enforce the target noise level

Intuition: add-then-remove

Concrete example

Sampled clients $|S| = 4$

Minimum necessary noise level $\sigma_*^2 = 1$

Dropout-Resilient Noise Enforcement

Goal: always enforce the target noise level

Intuition: add-then-remove

Concrete example

Add

Sampled clients $|S| = 4$

Dropout tolerance $t = 2$,

Minimum necessary noise level $\sigma_*^2 = 1$

Each client adds noise $n_i \sim \chi(1/2)$
to tolerate up to 2 clients to drop

Dropout-Resilient Noise Enforcement

Goal: always enforce the target noise level

Intuition: add-then-remove

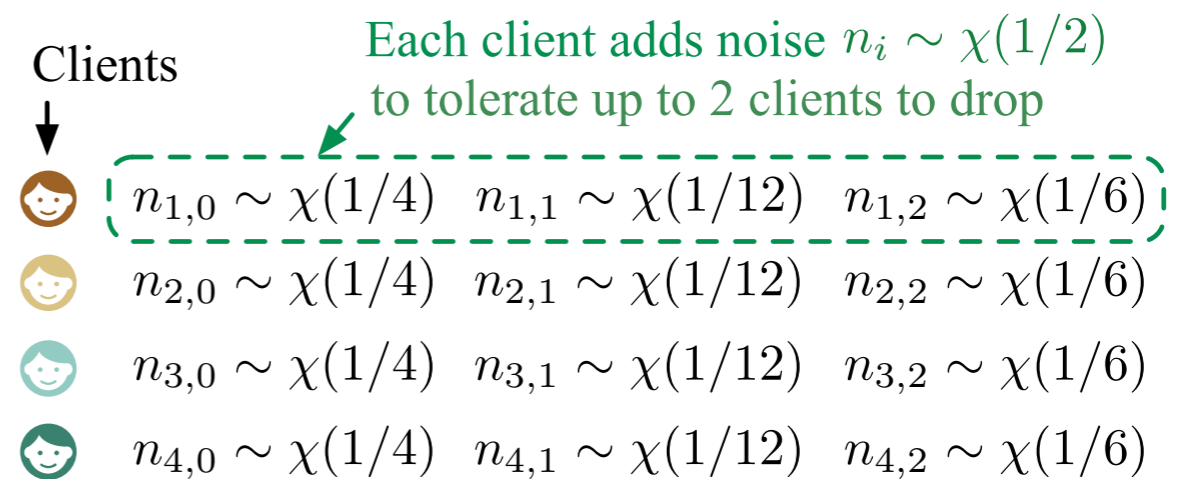
Concrete example

Add

Sampled clients $|S| = 4$

Dropout tolerance $t = 2$,

Minimum necessary noise level $\sigma_*^2 = 1$



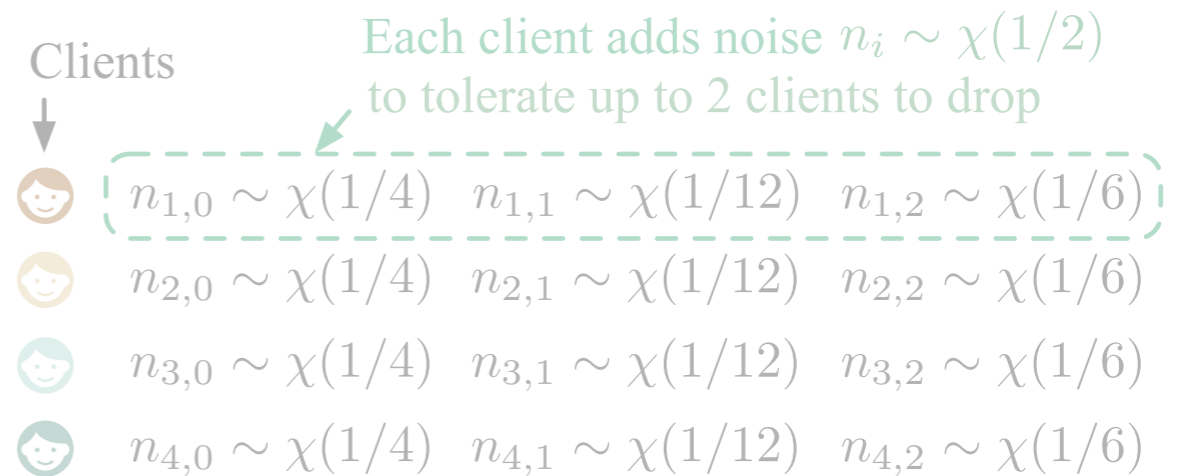
Dropout-Resilient Noise Enforcement

Goal: always enforce the target noise level

Intuition: add-then-remove

Concrete example


Add
 Sampled clients $|S| = 4$
 Dropout tolerance $t = 2$,
 Minimum necessary noise level $\sigma_*^2 = 1$



If 0 client drops

Achieve target noise $\sigma_*^2 = 1$

Remove



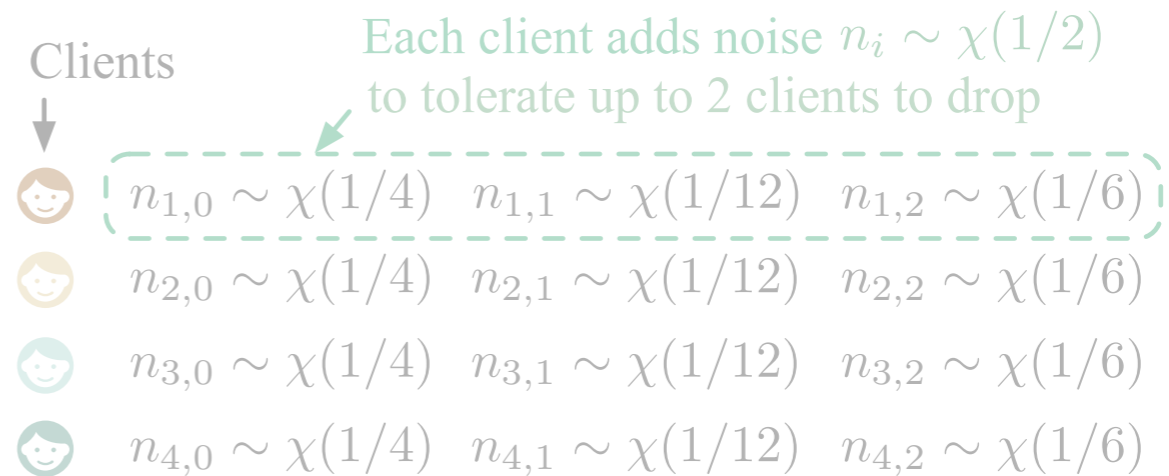
Dropout-Resilient Noise Enforcement

Goal: always enforce the target noise level


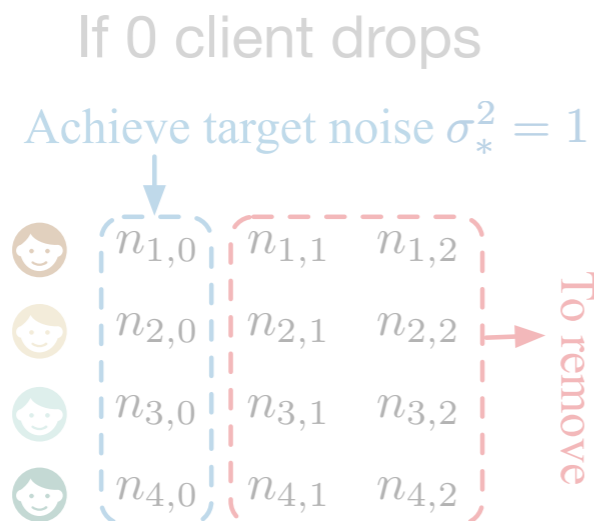
Intuition: add-then-remove

Concrete example

Add
 Sampled clients $|S| = 4$
 Dropout tolerance $t = 2$,
 Minimum necessary noise level $\sigma_*^2 = 1$



Remove




Dropout-Resilient Noise Enforcement

Goal: always enforce the target noise level

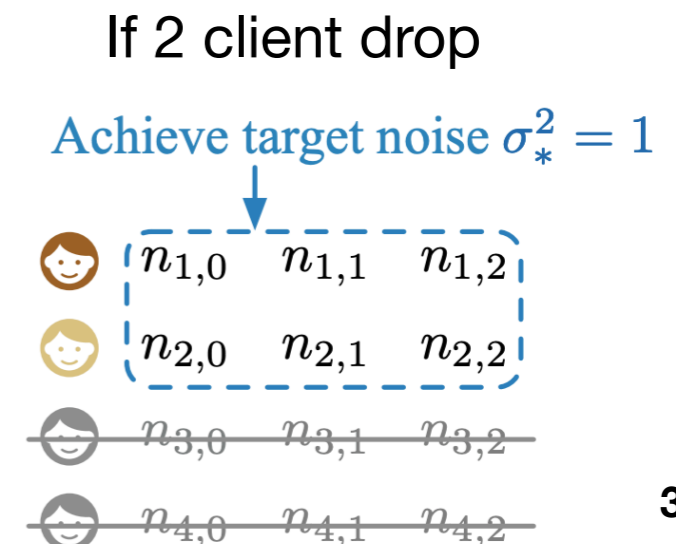

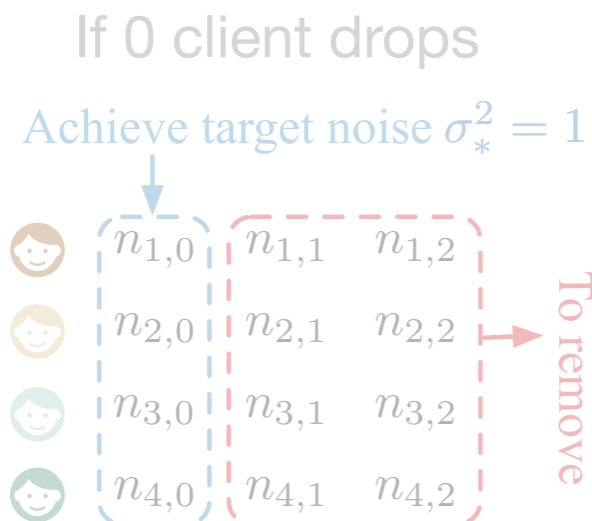
Intuition: add-then-remove

Concrete example

Add
 Sampled clients $|S| = 4$
 Dropout tolerance $t = 2$,
 Minimum necessary noise level $\sigma_*^2 = 1$



Remove



Dropout-Resilient Noise Enforcement

Goal: always enforce the target noise level

Intuition: add-then-remove

Concrete example

Formal definition: **XNoise**

- **Add**: decompose i 's added noise $n_i \sim \chi\left(\frac{\sigma_*^2}{|S| - t}\right)$ into $t + 1$ components:

$$n_i = \sum_{k=0}^t n_{i,k}, n_{i,0} \sim \chi\left(\frac{\sigma_*^2}{|S|}\right), \text{ and } n_{i,k} \sim \chi\left(\frac{\sigma_*^2}{(|S| - k + 1)(|S| - k)}\right) (k \in [t])$$

- **Remove**: when there are $|D|$ clients dropping out, the noise components $n_{i,k}$ contributed by the surviving clients $i \in S \setminus D$ with the index $k > |D|$ becomes excessive and is removed by the server

Practical Design:

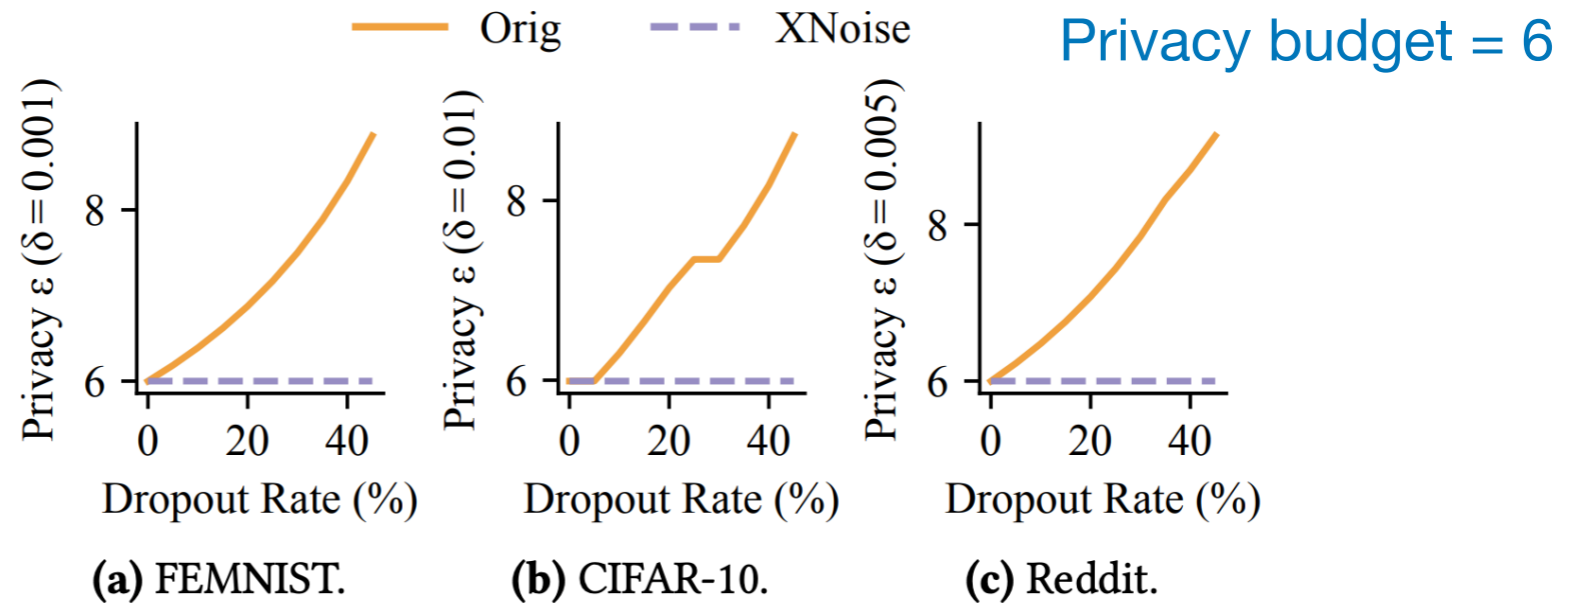
- **Avoiding cascading dropout**: secret sharing
- **Integrity of the dropout outcome**: secure signature

Dropout-Resilient Noise Enforcement

XNoise

Improves privacy

without sacrificing
final model utility




Dropout rates

<i>d</i>	0		10%		20%		30%		40%	
	Ori	XNo	Ori	XNo	Ori	XNo	Ori	XNo	Ori	XNo
F	61.3	61.4	61.4	61.4	61.2	61.4	61.2	61.2	61.4	61.5
C	66.5	66.3	66.7	66.9	66.6	65.7	64.3	65.7	63.8	64.2
R	2169	2142	2158	2179	2286	2285	2294	2317	2299	2329

Dropout-Resilient Noise Enforcement

XNoise

Improves privacy




without sacrificing
final model utility

Dropout rates

Datasets	0		10%		20%		30%		40%	
	Ori	XNo	Ori	XNo	Ori	XNo	Ori	XNo	Ori	XNo
F	61.3	61.4	61.4	61.4	61.2	61.4	61.2	61.2	61.4	61.5
C	66.5	66.3	66.7	66.9	66.6	65.7	64.3	65.7	63.8	64.2
R	2169	2142	2158	2179	2286	2285	2294	2317	2299	2329

and incurs
acceptable
($\leq 34\%$)

runtime cost



Example: no dropout

Distributed DP has Two **Issues**

1. **Privacy** Issue: caused by client dropout
2. **Performance** Issue: **expensive** use of secure aggregation

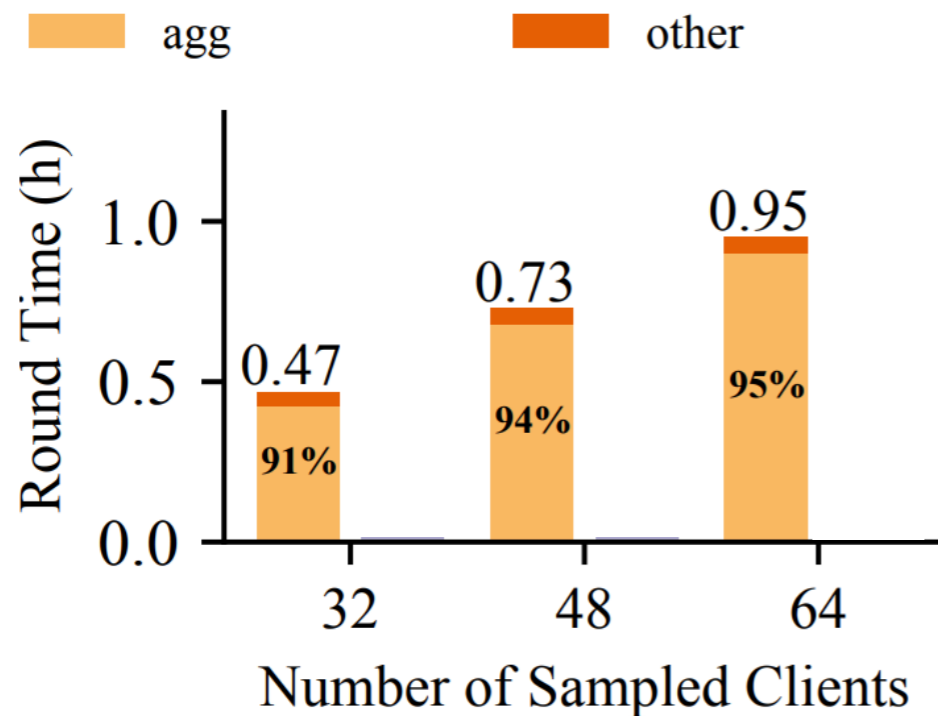
Performance issues with SecAgg

Extensive use of [secret sharing](#) and [pairwise masking](#)

Performance issues with SecAgg

Extensive use of secret sharing and pairwise masking

Dominates the training time (at least 91%)



original secure aggregation:

SecAgg


Performance issues with SecAgg

Extensive use of secret sharing and pairwise masking

Dominates the training time (at least 91%)

Follow-up solutions

- e.g. [SecAgg+](#): improves asymptotically



original secure aggregation:

SecAgg


Performance issues with SecAgg

Extensive use of secret sharing and pairwise masking


Dominates the training time (at least 91%)

Follow-up solutions have **inefficiencies**

- e.g. SecAgg+: improves asymptotically, but **help little** in small-scale practice¹



original secure aggregation:
SecAgg



SOTA secure aggregation:
SecAgg+

Pipeline-Parallel Acceleration

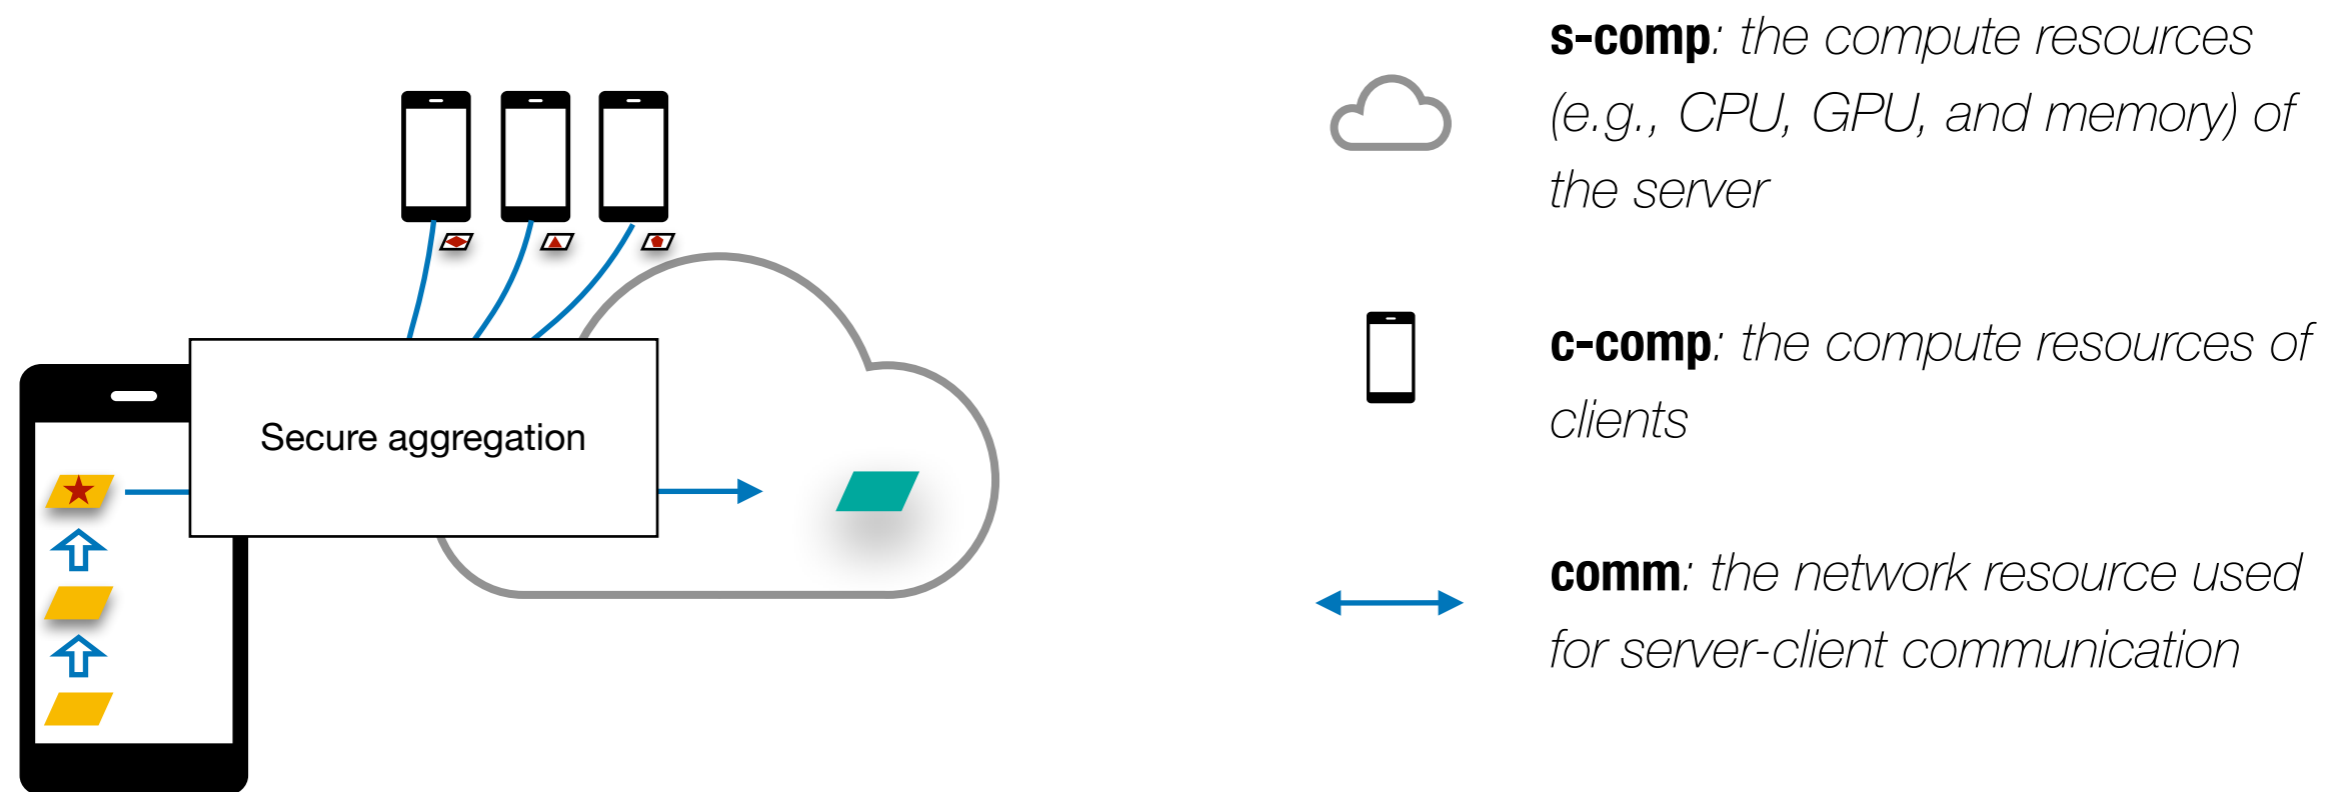
Goal: leverage the **underutilized resources** in the system level

Pipeline-Parallel Acceleration

Goal: leverage the underutilized resources in the system level

Approach:

- Step 1: Identify the **types** of system resources




Pipeline-Parallel Acceleration

Goal: leverage the underutilized resources in the system level

Approach:

- Step 1: Identify the types of system resources
- Step 2: **Group** consecutive operations that use the same system resources




Step	Operation	Stage (Resource)
1	Clients encode updates.	
2	Clients generate security keys.	
3	Clients establish shared secrets.	1 (c-comp)
4	Clients mask encoded updates.	
5	Clients upload masked updates.	2 (comm)
6	Server deals with dropout.	
7	Server computes aggregate update.	3 (s-comp)
8	Server updates the global model.	
9	Server dispatches the aggregate.	4 (comm)
10	Clients decode the aggregate.	
11	Clients use the aggregate.	5 (c-comp)

Pipeline-Parallel Acceleration

Goal: leverage the underutilized resources in the system level

Approach:

- Step 1: Identify the types of system resources
- Step 2: Group consecutive operations that use the same system resources
- Step 3: Evenly **partition** each client's update into chunks and **pipeline** them



Pipeline-Parallel Acceleration

Goal: leverage the underutilized resources in the system level

Approach:

- Step 1: Identify the types of system resources
- Step 2: Group consecutive operations that use the same system resources
- Step 3: Evenly partition each client's update into chunks and pipeline them
 - Optimize to determine the **optimal number of chunk, m^***

$$m^* = \arg \min_{m \in \mathbb{N}_+} f_{a,m}$$

$$s.t. \quad f_{s,c} = b_{s,c} + l_s$$

$$b_{s,c} = \max\{o_{s,c}, r_{s,c}\}$$

$$o_{s,c} = \begin{cases} 0, & \text{if } s = 0, \\ f_{s-1,c} & \text{Intra-chunk sequential execution} \end{cases}$$

*Definition of the finish time of
Chunk m at Stage a*

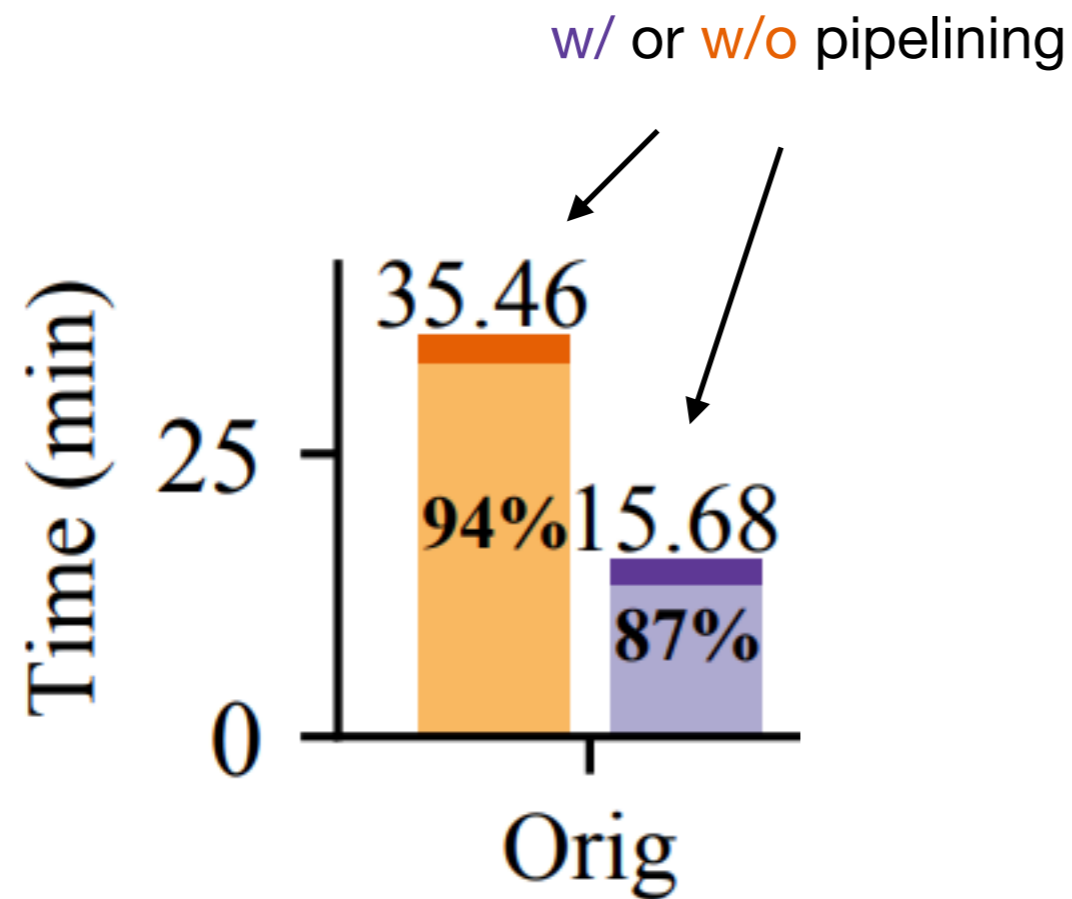
$$r_{s,c} = \begin{cases} 0, & \text{if } s = 0 \text{ and } c = 0, \\ f_{q,m} \text{ or } \perp, & \text{if } s \neq 0 \text{ and } c = 0, \\ f_{s,c-1}, & \text{otherwise} \end{cases}$$

*Exclusive allocation
& Inter-chunk sequential execution*

Pipeline-Parallel Acceleration

Effectiveness:

① A maximum speedup of 2.4×





Case study: CIFAR10 @

VGG19, dropout rate = 30%

Pipeline-Parallel Acceleration

Effectiveness:

① A maximum speedup of 2.4×




w/o or w/
our noise enforcement

Pipeline-Parallel Acceleration

Effectiveness:

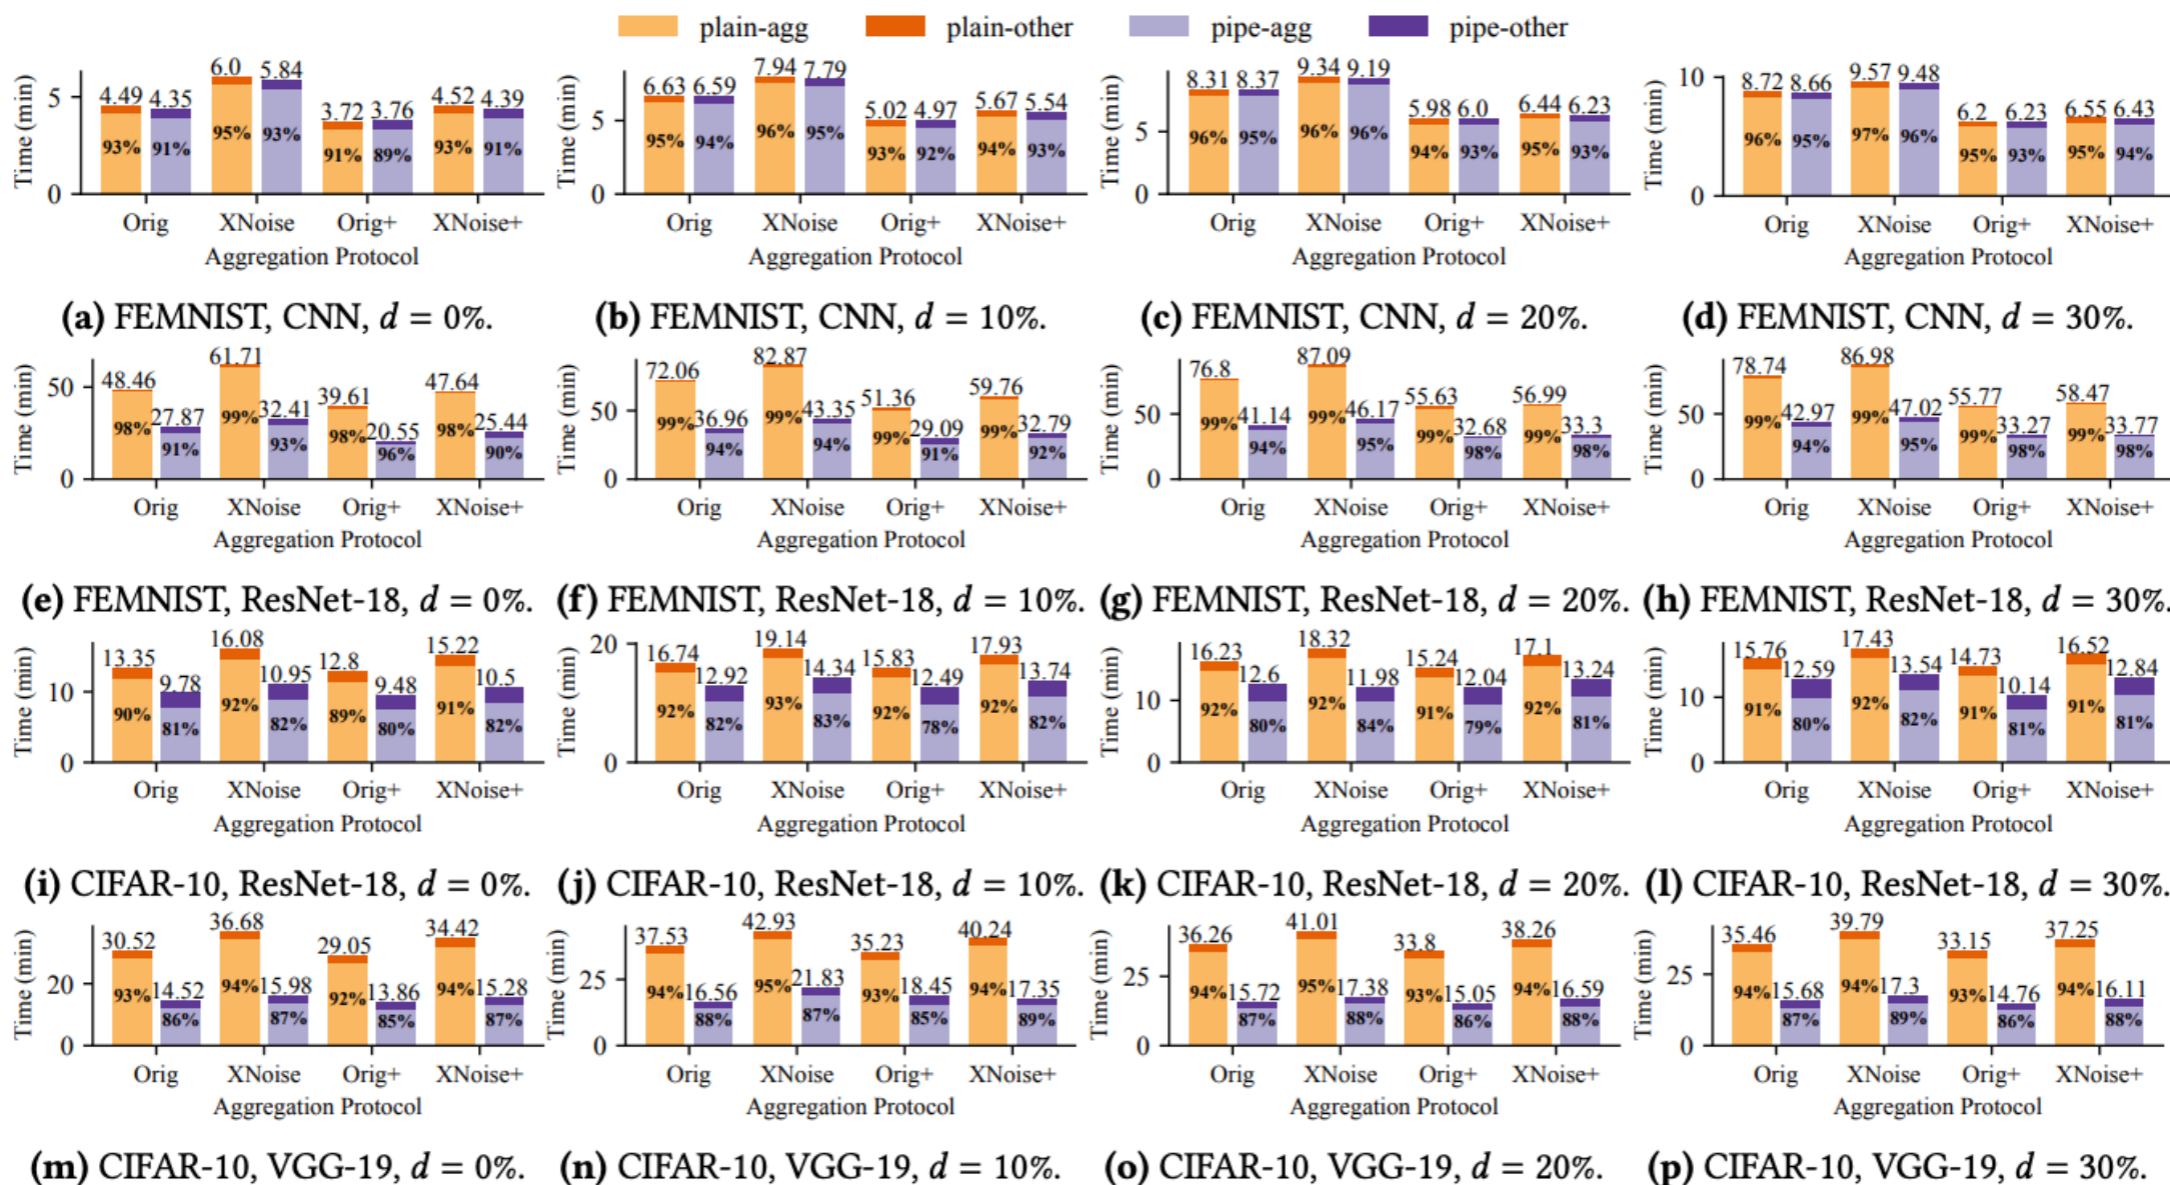
① A maximum speedup of 2.4x



Pipeline-Parallel Acceleration

Effectiveness:

① A maximum speedup of 2.4X



②

Larger models


gain more

③

Scaling w/ # part.

④ Gains are consistent across different dropout rates

Dordis

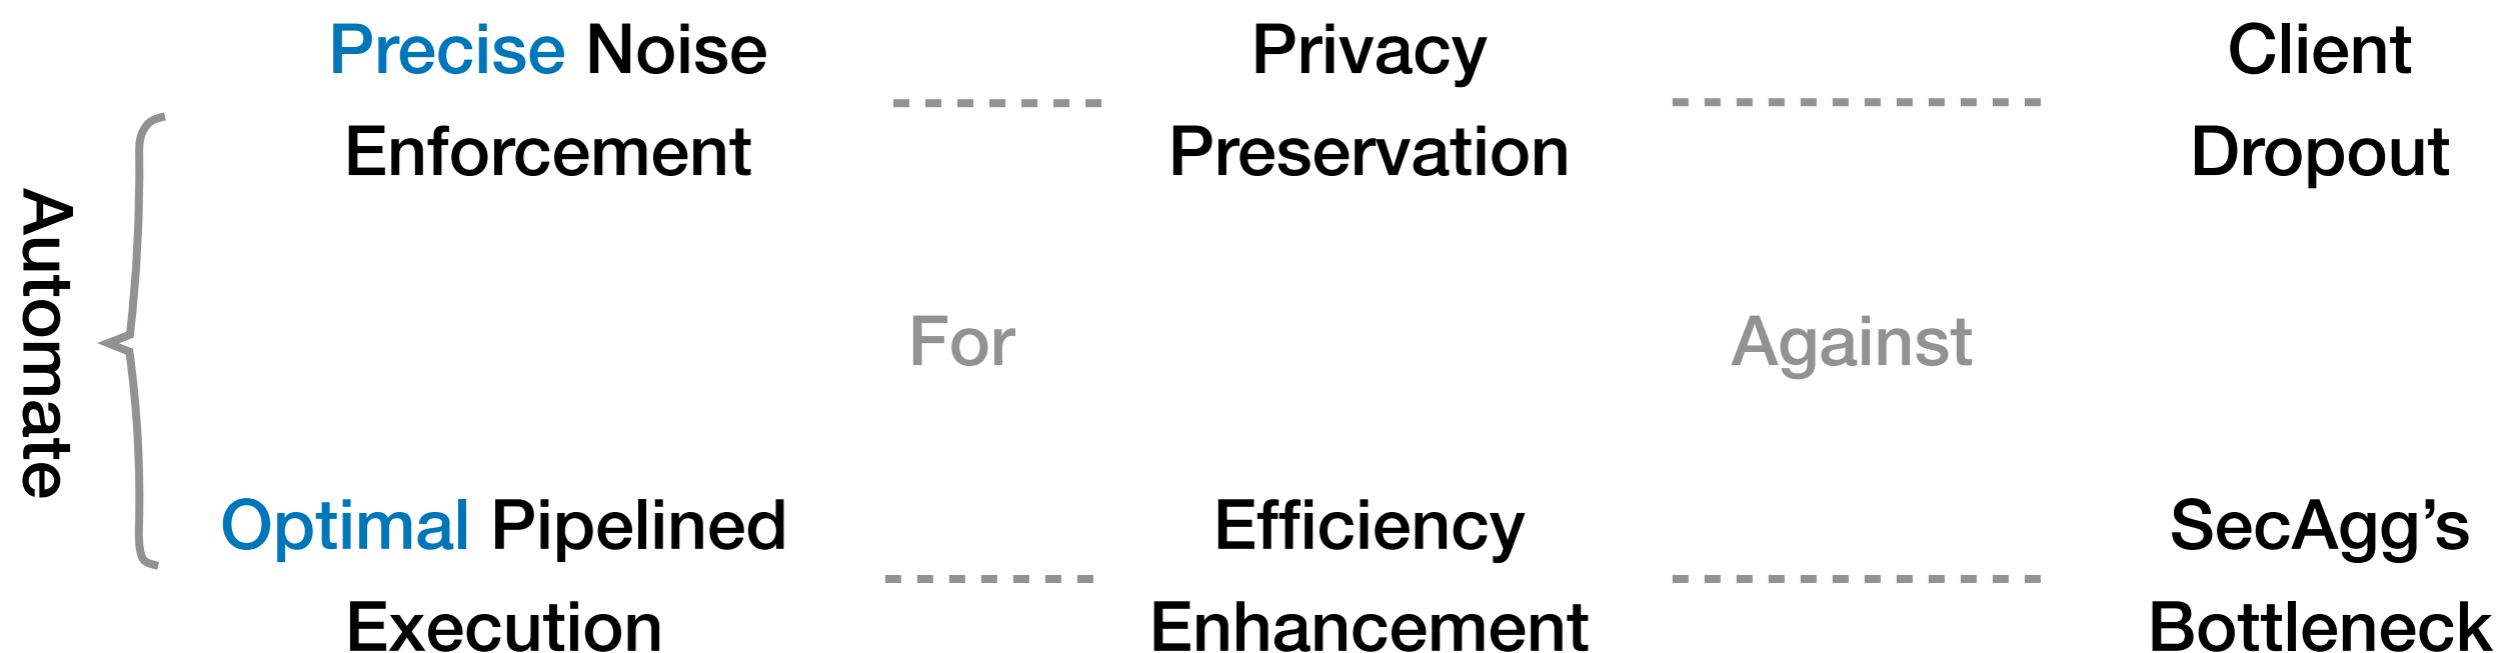


A distributed DP framework for

- Privacy
- Efficiency

in FL training

<https://github.com/SamuelGong/Dordis>



Thank you!