

Securing Secure Aggregation: Mitigating Multi-Round Privacy Leakage in Federated Learning

Ramy E. Ali



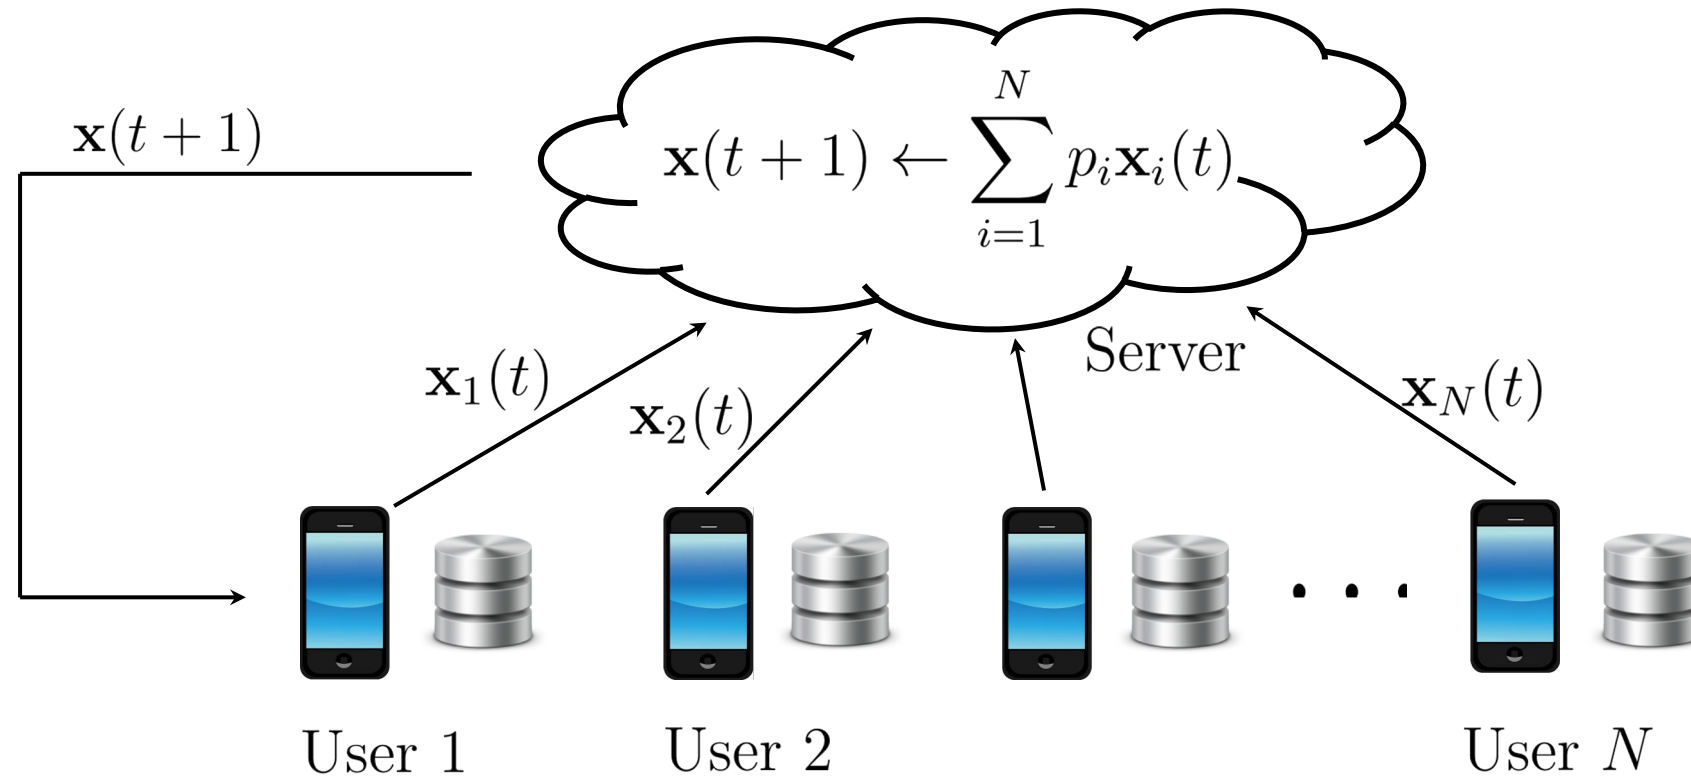
USC University of
Southern California

In collaboration with Jinhyun So (USC), Başak Güler (UCR), Salman Avestimehr (USC) and Jiantao Jiao (UC Berkeley)

2021

The Promise of Federated Learning

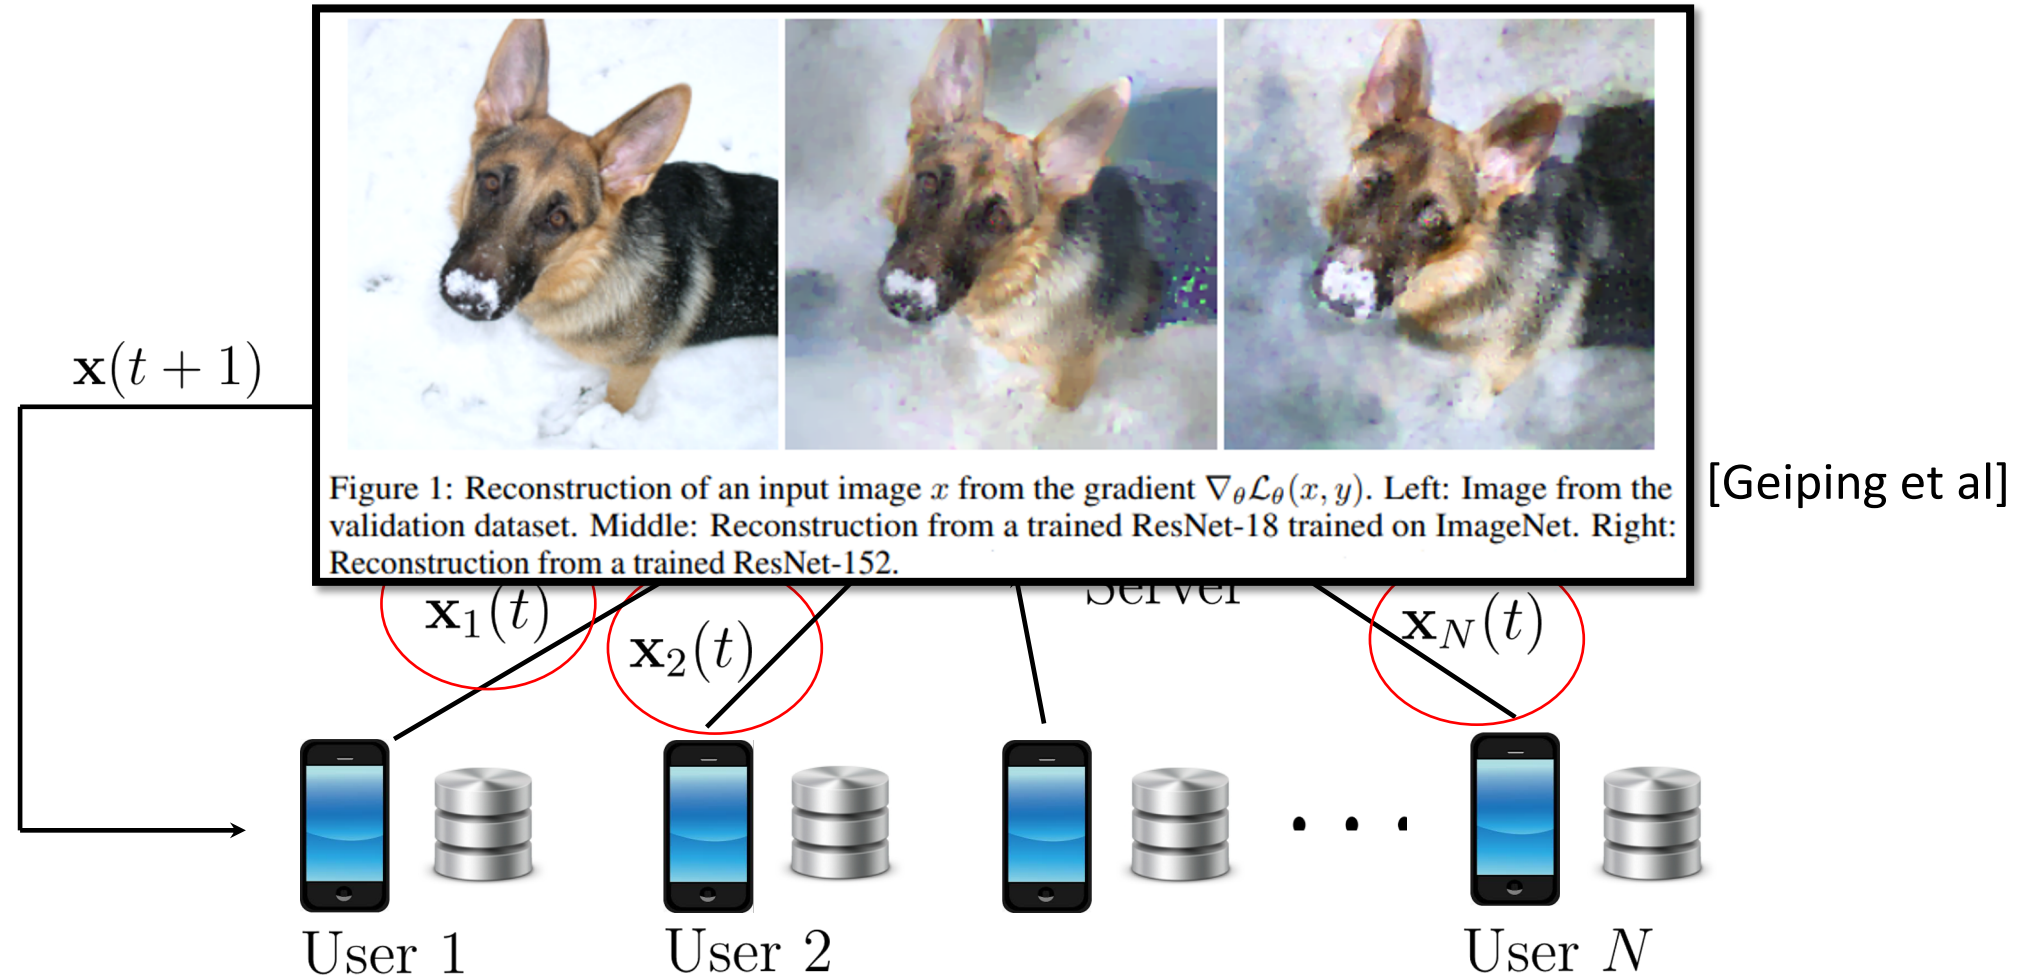
Main principle: train locally - average globally



Ensuring **privacy** by avoiding data sharing?

But ...

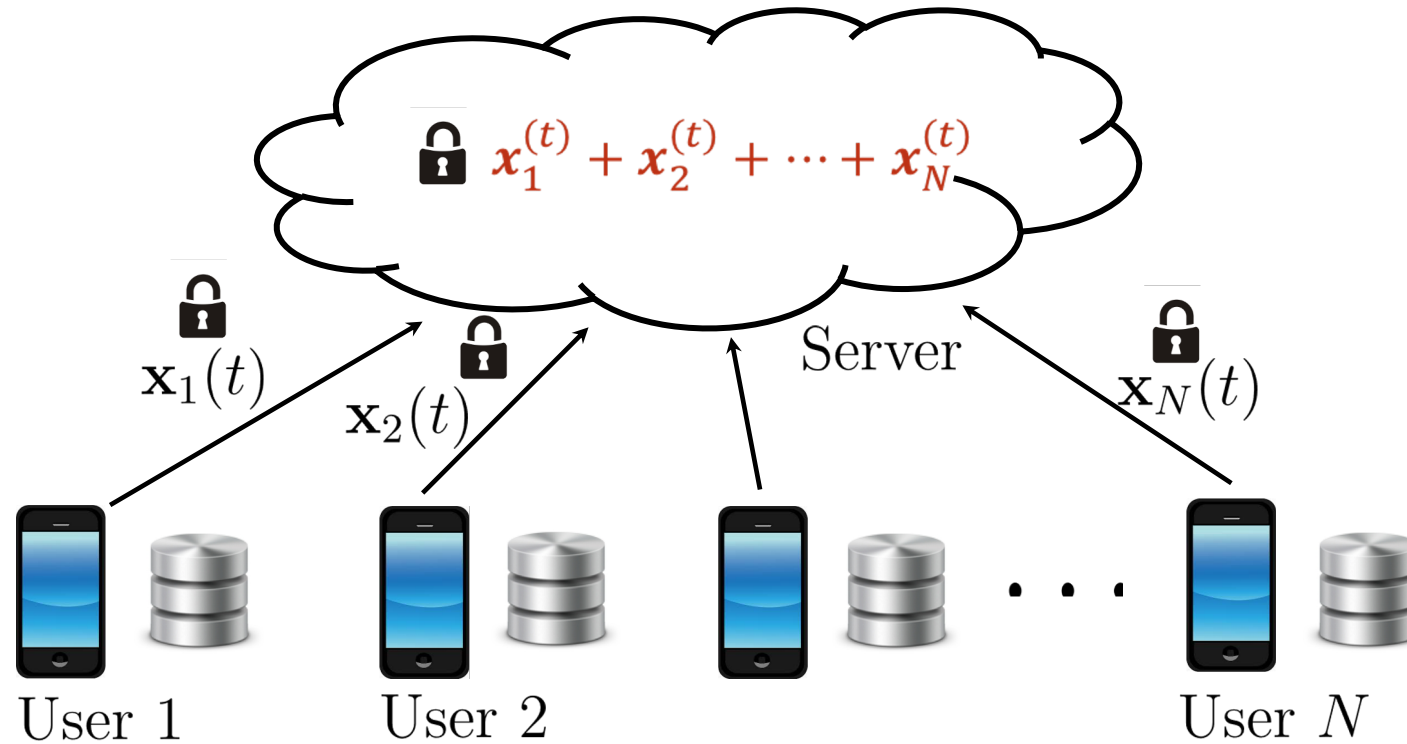
Model Inversion Attack



Problem: Individual model updates can leak sensitive data

Remedy: Secure Model Aggregation

- Secure aggregation ensures that the server only learns the global model.



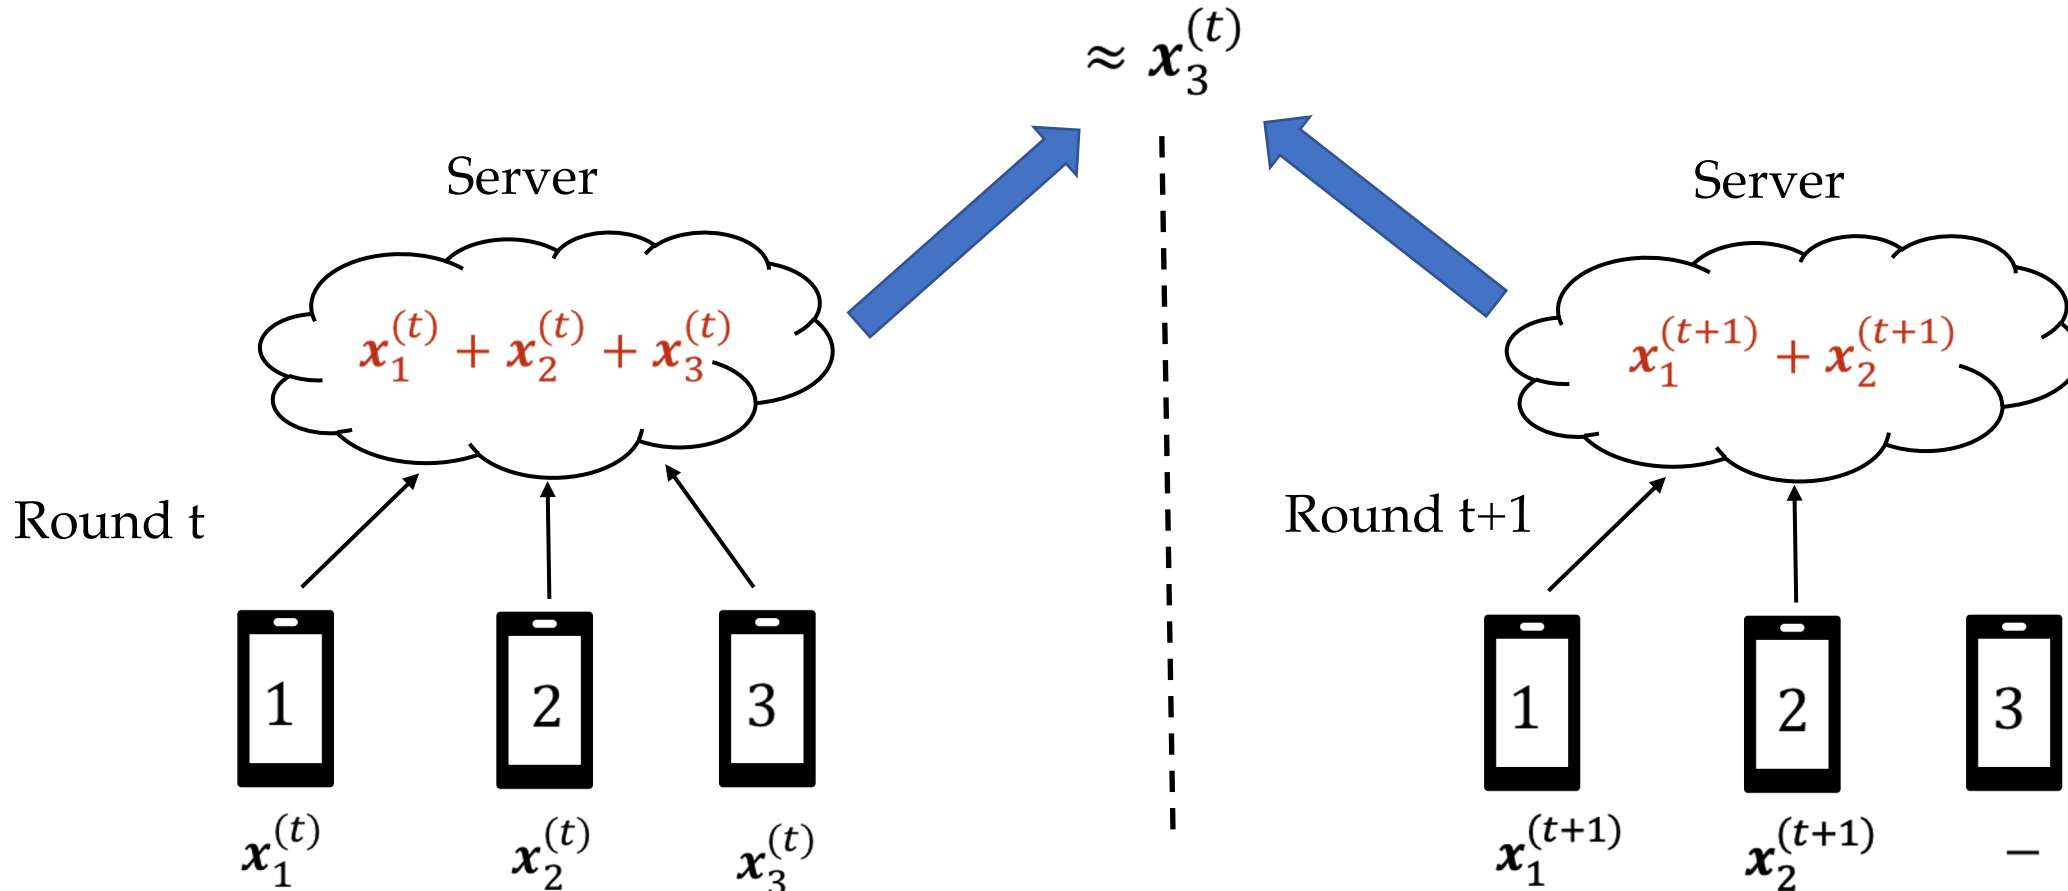
Secure Aggregation is Essentially an MPC problem with [User Dropouts](#)

Bad news ...

- Secure aggregation, however, is not secure over **multiple** rounds.

Bad news ...

- Secure aggregation, however, is not secure over **multiple** rounds.
- **Intuition:** partial user participation leads to privacy leakage




This is a serious issue!

- Random selection may reveal all individual models.
- **Experiment**
 - N=40 users
 - MNIST dataset with non-IID distribution
 - K=8 users are selected at random at each round
 - The server estimates the individual gradients through least-squares

Reconstruction error of user i at time t

$$\frac{\|\mathbf{x}_i^{(t)} - \hat{\mathbf{x}}_i^{(t)}\|^2}{\|\mathbf{x}_i^{(t)}\|^2}$$



reconstruction error after 600 rounds


reconstruction error < 0.005 (for many users)

This is a serious issue!

- Random selection may reveal all individual models.
- **Experiment**
 - N=40 users
 - MNIST dataset with non-IID distribution
 - K=8 users are selected at random at each round
 - The server estimates the individual gradients through least-squares

Reconstruction error of user i at time t

$$\frac{\|\mathbf{x}_i^{(t)} - \hat{\mathbf{x}}_i^{(t)}\|^2}{\|\mathbf{x}_i^{(t)}\|^2}$$



reconstruction error < 0.005 (for many users)

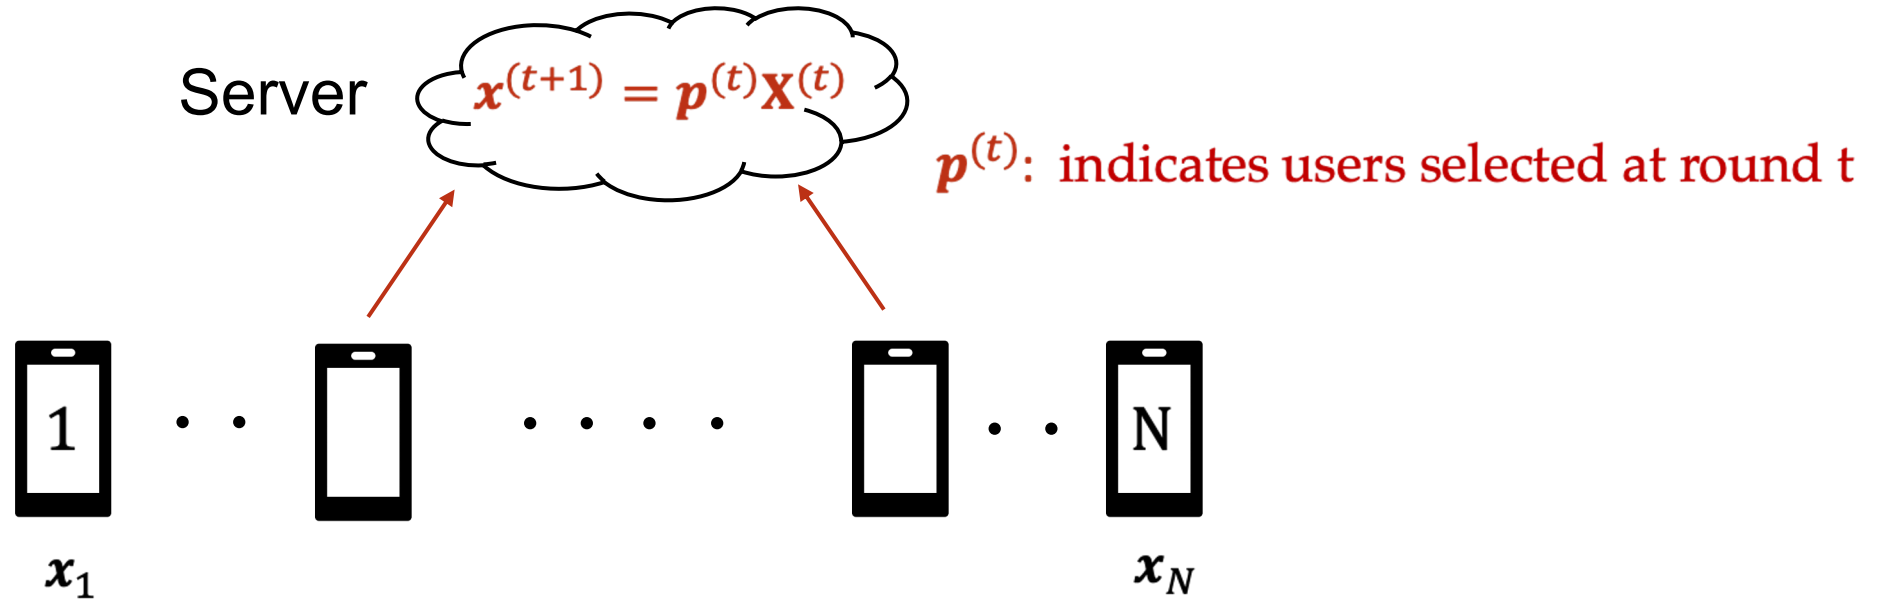
Similar concerns have also been reported in some recent works:

- Pejó et al. "Quality Inference in Federated Learning with Secure Aggregation.", 2020.
- M. Lam et al. "Gradient Disaggregation: Breaking Privacy in Federated Learning by Reconstructing the User Participant Matrix.", ICML 2021.

This talk

- Introduce a notion/metric for **multi-round privacy**
- Propose **Multi-RoundSecAgg**, which ensures multi-round privacy
 - It further optimizes the average number of participating users (convergence rate) and fairness in user selection
 - It also introduces a trade-off between “privacy” and “convergence rate”

Federated Averaging with Partial User Participation

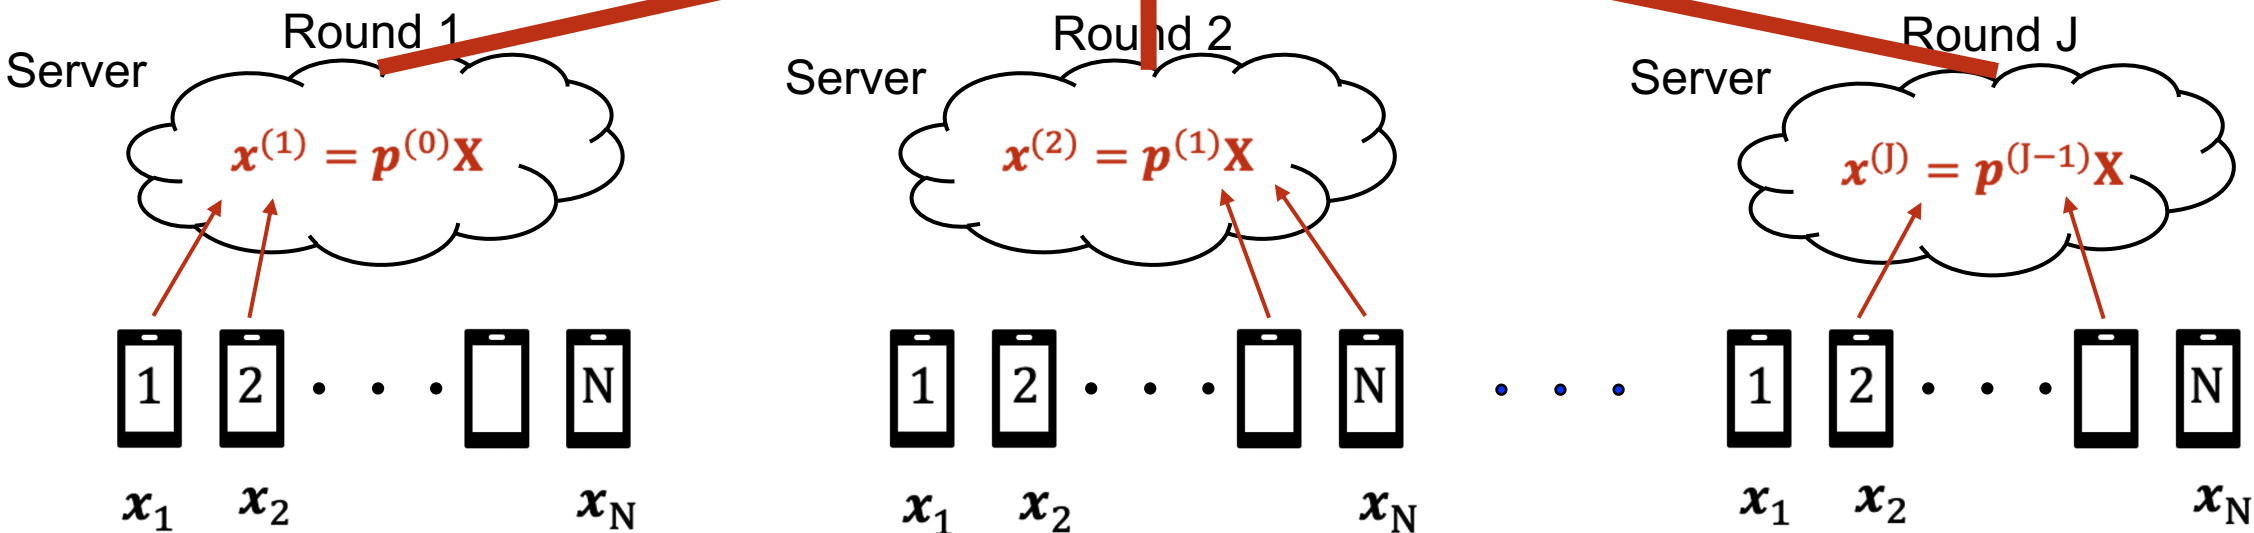


- Participation matrix $\mathbf{P}^{(J)} = \begin{pmatrix} \mathbf{p}^{(0)} \\ \vdots \\ \mathbf{p}^{(J-1)} \end{pmatrix} \in \{0,1\}^{J \times N}$, J : number of rounds

How should we choose $\mathbf{P}^{(J)}$ to ensure long-term privacy?

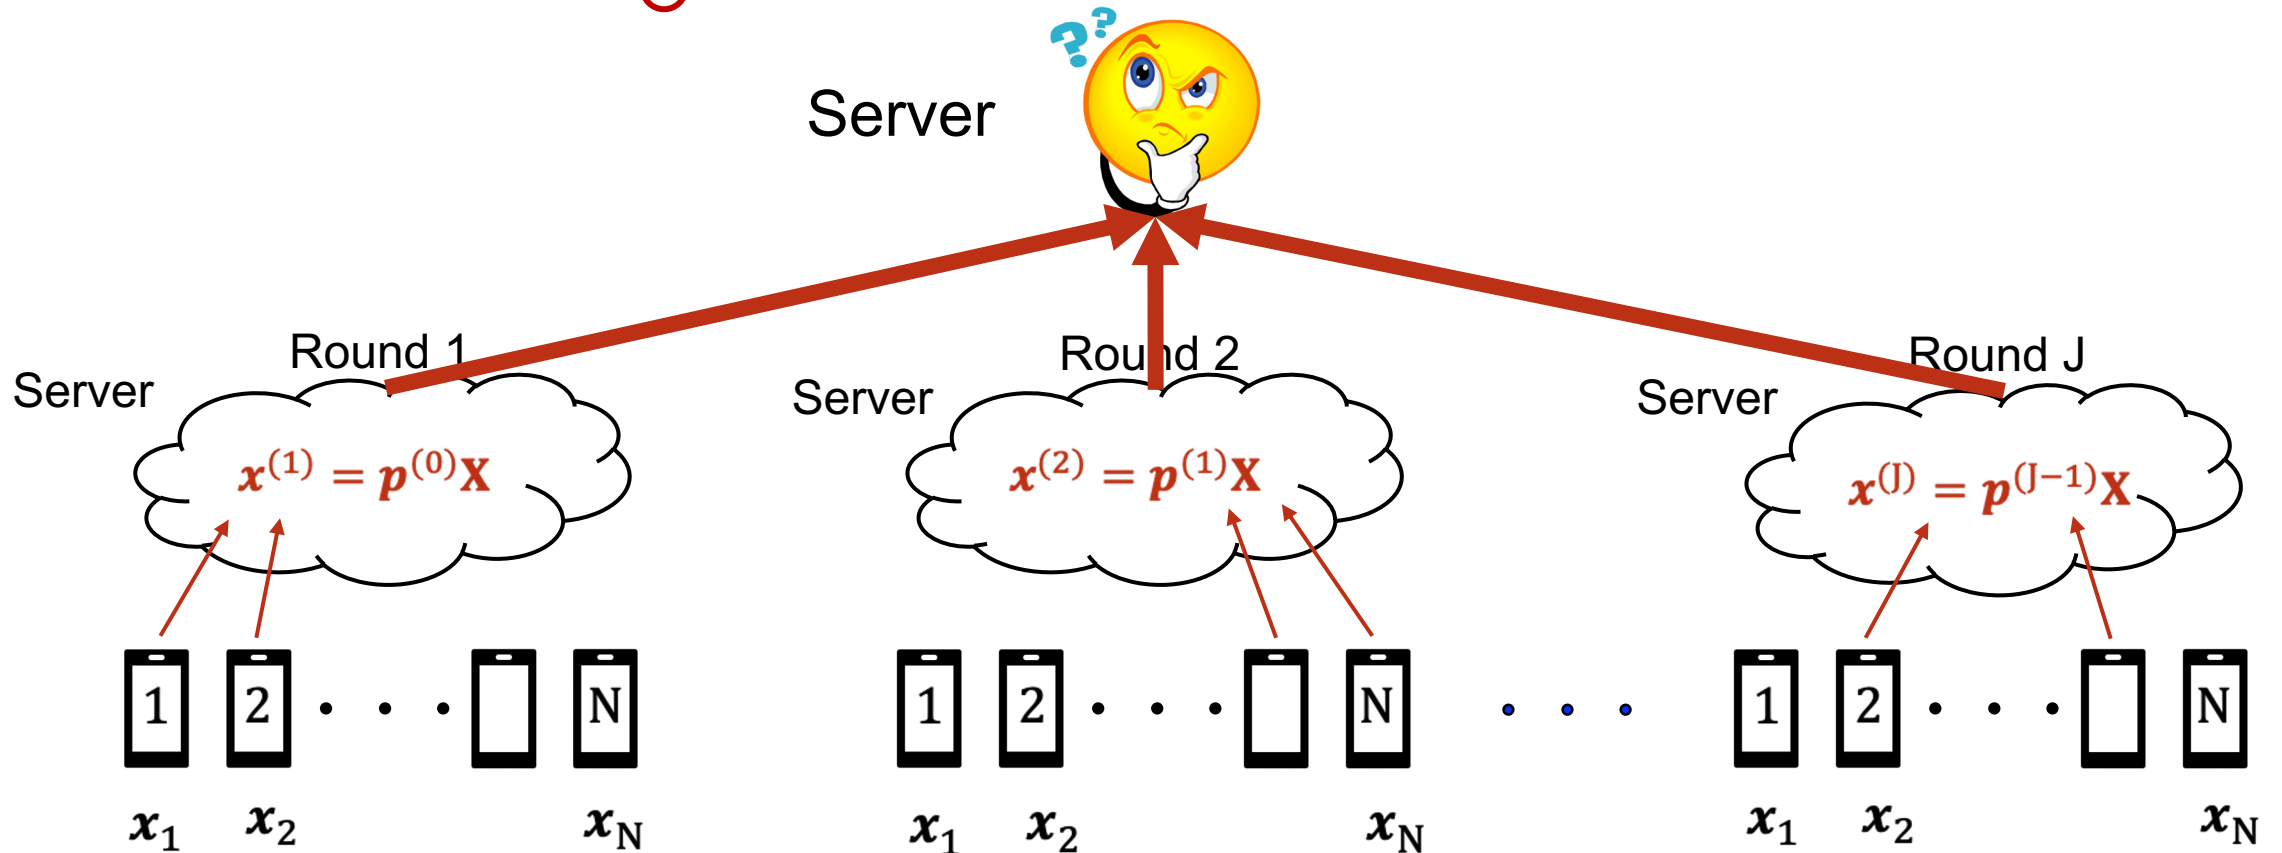
Metric 1: Multi-round Privacy (T)

7801	2229	2226	2263	2307	2307	2361	2381	2401	2411	2422	2469	1	24	55	83	100	136	158	196	237	240	257	266	326	337	364	411	454	465	475	478	484	485	487	488	489	490	491	492	493	494	495	496	497	498	499	500	501	502	503	504	505	506	507	508	509	510	511	512	513	514	515	516	517	518	519	520	521	522	523	524	525	526	527	528	529	530	531	532	533	534	535	536	537	538	539	540	541	542	543	544	545	546	547	548	549	550	551	552	553	554	555	556	557	558	559	560	561	562	563	564	565	566	567	568	569	570	571	572	573	574	575	576	577	578	579	580	581	582	583	584	585	586	587	588	589	590	591	592	593	594	595	596	597	598	599	600	601	602	603	604	605	606	607	608	609	610	611	612	613	614	615	616	617	618	619	620	621	622	623	624	625	626	627	628	629	630	631	632	633	634	635	636	637	638	639	640	641	642	643	644	645	646	647	648	649	650	651	652	653	654	655	656	657	658	659	660	661	662	663	664	665	666	667	668	669	670	671	672	673	674	675	676	677	678	679	680	681	682	683	684	685	686	687	688	689	690	691	692	693	694	695	696	697	698	699	700
------	------	------	------	------	------	------	------	------	------	------	------	---	----	----	----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----



Metric 1: Multi-round Privacy (T)

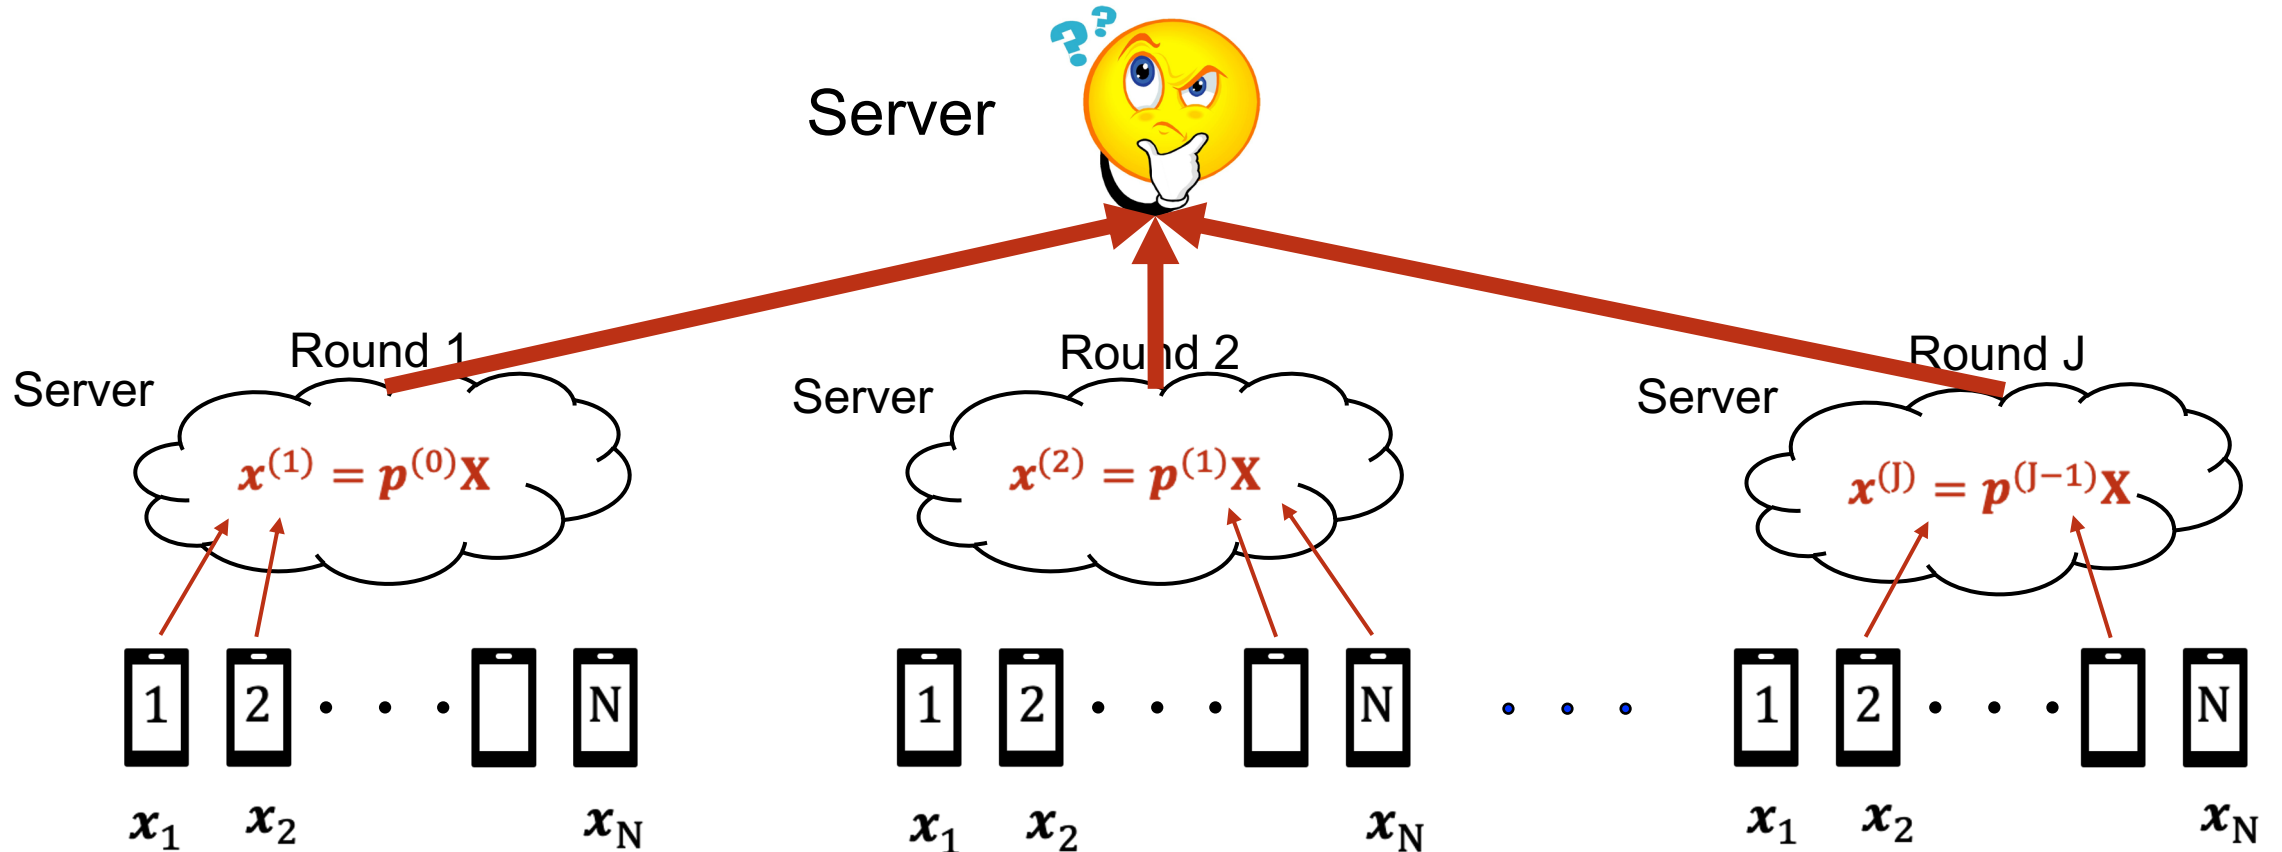
The server must not learn any aggregate model of less than T users.



Metric 1: Multi-round Privacy (T)

- A multi-round privacy T requires that any non-zero partial sum that the server can reconstruct to be of the form

$$a_1 \sum_{j \in \mathcal{S}_1} \mathbf{x}_j + a_2 \sum_{j \in \mathcal{S}_2} \mathbf{x}_j + \dots + a_n \sum_{j \in \mathcal{S}_n} \mathbf{x}_j, \text{ where } |\mathcal{S}_i| \geq T.$$



Metric 1: Multi-round Privacy (T)

- A multi-round privacy T requires that any non-zero partial sum that the server can reconstruct to be of the form

$$a_1 \sum_{j \in \mathcal{S}_1} \mathbf{x}_j + a_2 \sum_{j \in \mathcal{S}_2} \mathbf{x}_j + \cdots + a_n \sum_{j \in \mathcal{S}_n} \mathbf{x}_j, \text{ where } |\mathcal{S}_i| \geq T.$$

- **Example** ($T=2$): the best the server can do is reconstructing x_i+x_j (for some i and j)

Metric 1: Multi-round Privacy (T)

- A multi-round privacy T requires that any non-zero partial sum that the server can reconstruct to be of the form

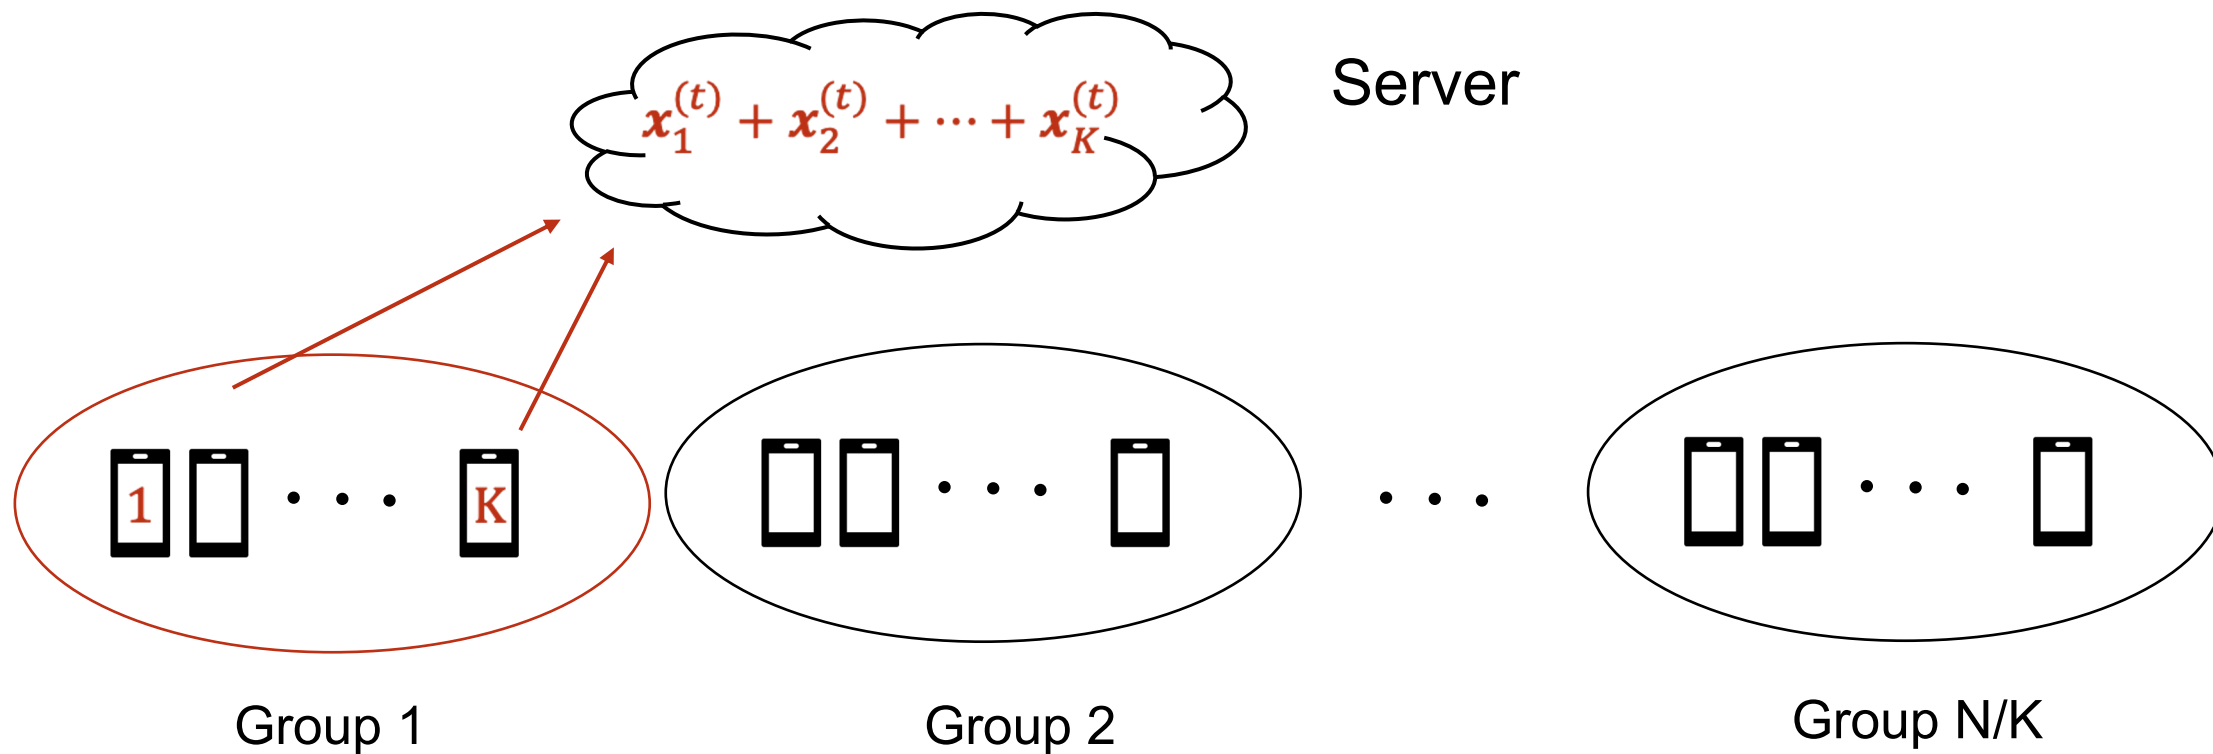
$$a_1 \sum_{j \in \mathcal{S}_1} \mathbf{x}_j + a_2 \sum_{j \in \mathcal{S}_2} \mathbf{x}_j + \dots + a_n \sum_{j \in \mathcal{S}_n} \mathbf{x}_j, \text{ where } |\mathcal{S}_i| \geq T.$$

- **Example** ($T=2$): the best the server can do is reconstructing x_i+x_j (for some i and j)
- Worst-case (strong) assumption 1: the model coefficients in each group are the same
- Worst-case (strong) assumption 2: user's model stays the same across different rounds

Baselines

1. User Partitioning

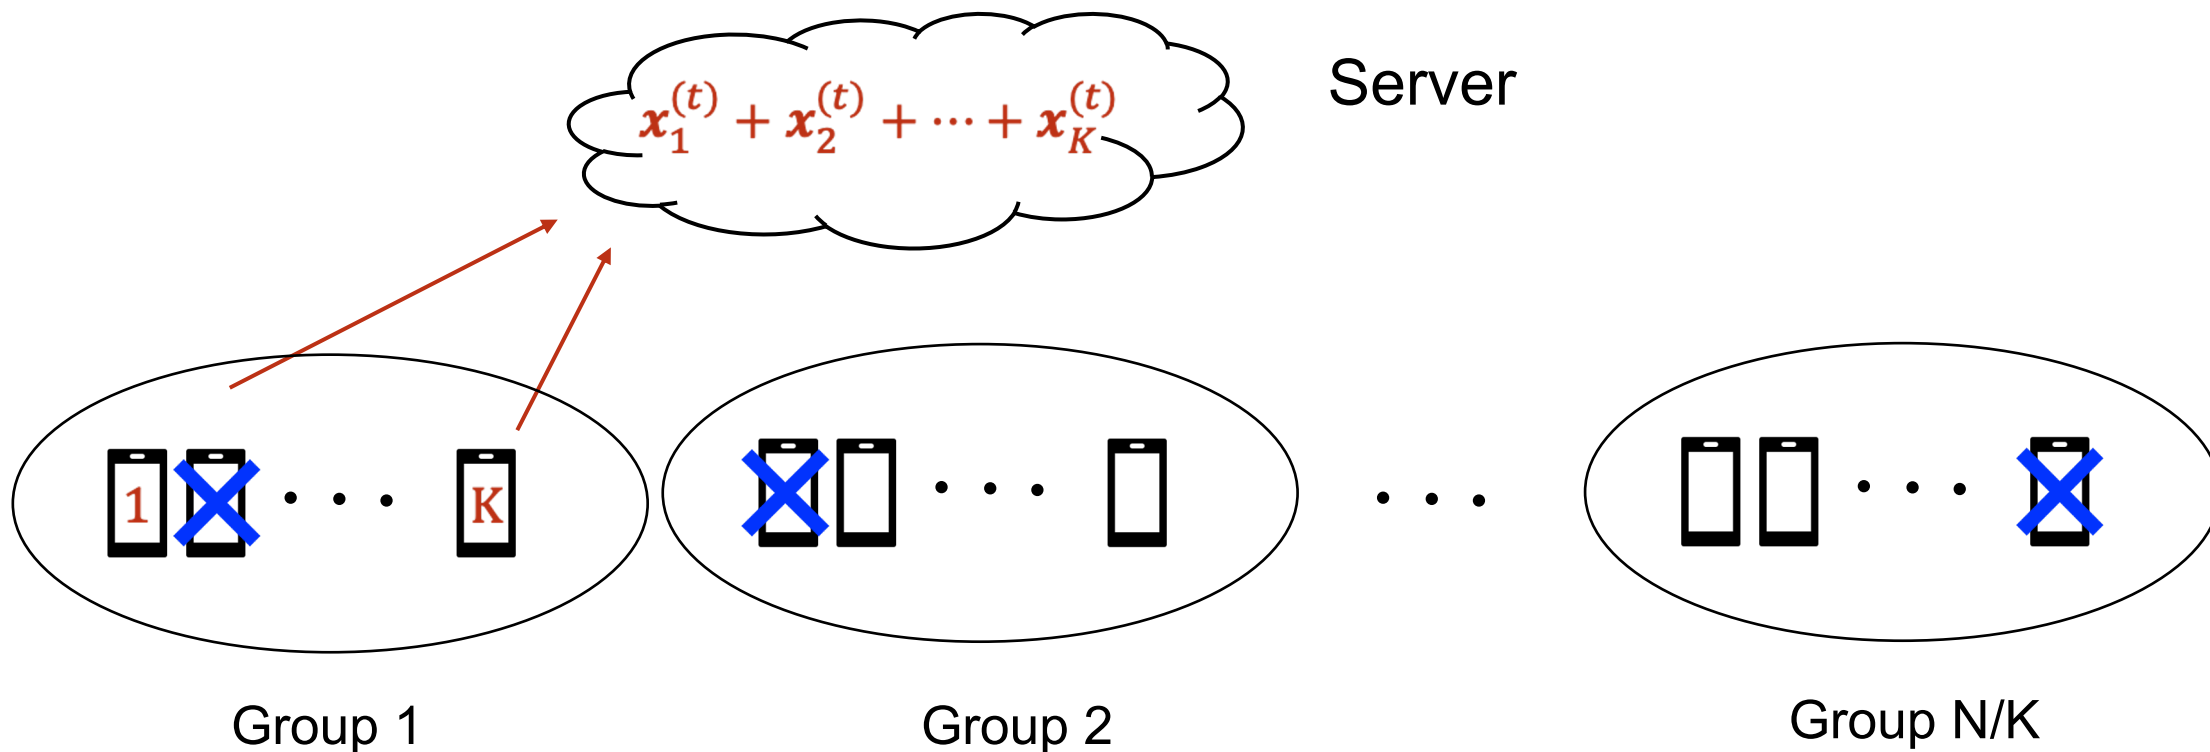
- Large multi-round privacy $T = \text{group size}$ 👍



Baselines

1. User Partitioning


- Large multi-round privacy $T = \text{group size}$ 👍
- In many rounds, however, no groups are available 👎



Baselines

2. Random Selection

- Small multi-round privacy $T = 1$ 🚫



Theorem: In Random Selection, the server can reconstruct all individual models of the N users after N rounds with probability at least

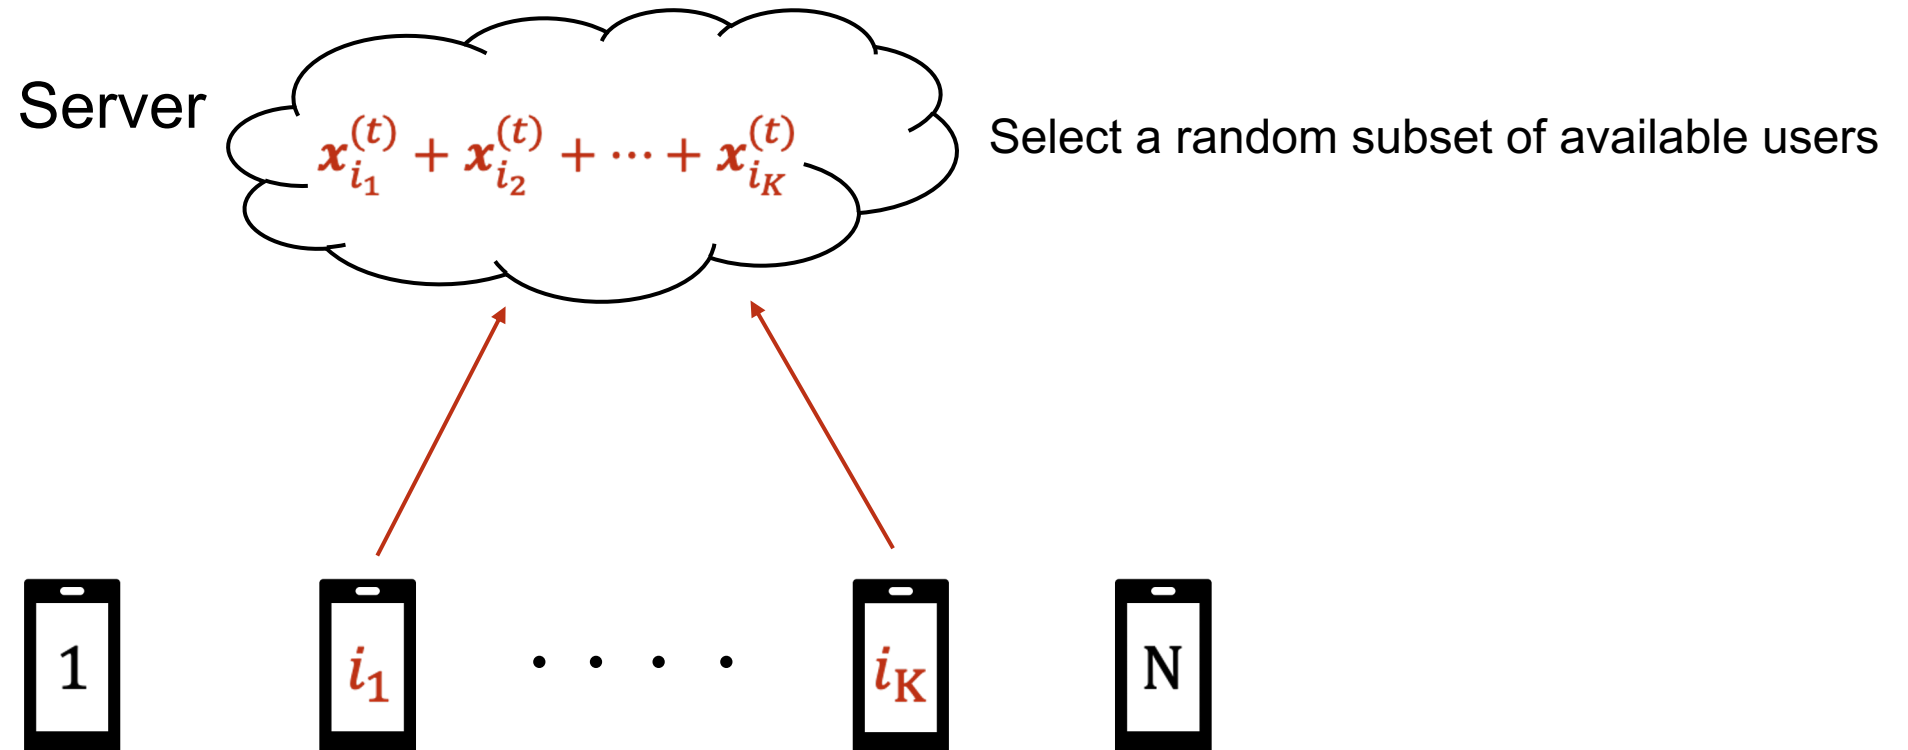
$$1 - \exp(-cN),$$

where c is a constant.

Baselines

2. Random Selection

- Small multi-round privacy $T = 1$ 🚫
- Any subset of available users can be selected in any round 👍



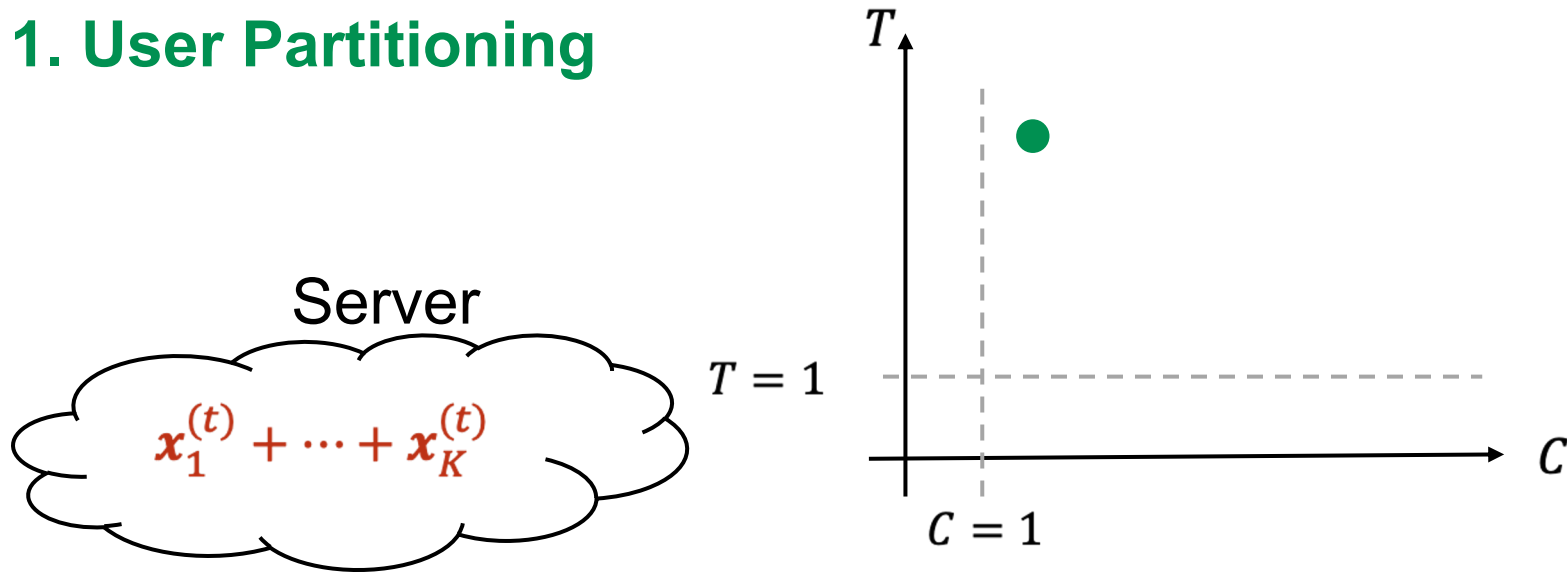
Metric 2: Average Aggregation Cardinality (C)

Aggregation Cardinality C = Average number of participating users over all rounds

Metric 2: Average Aggregation Cardinality (C)

Aggregation Cardinality C = Average number of participating users over all rounds

1. User Partitioning



Group 1

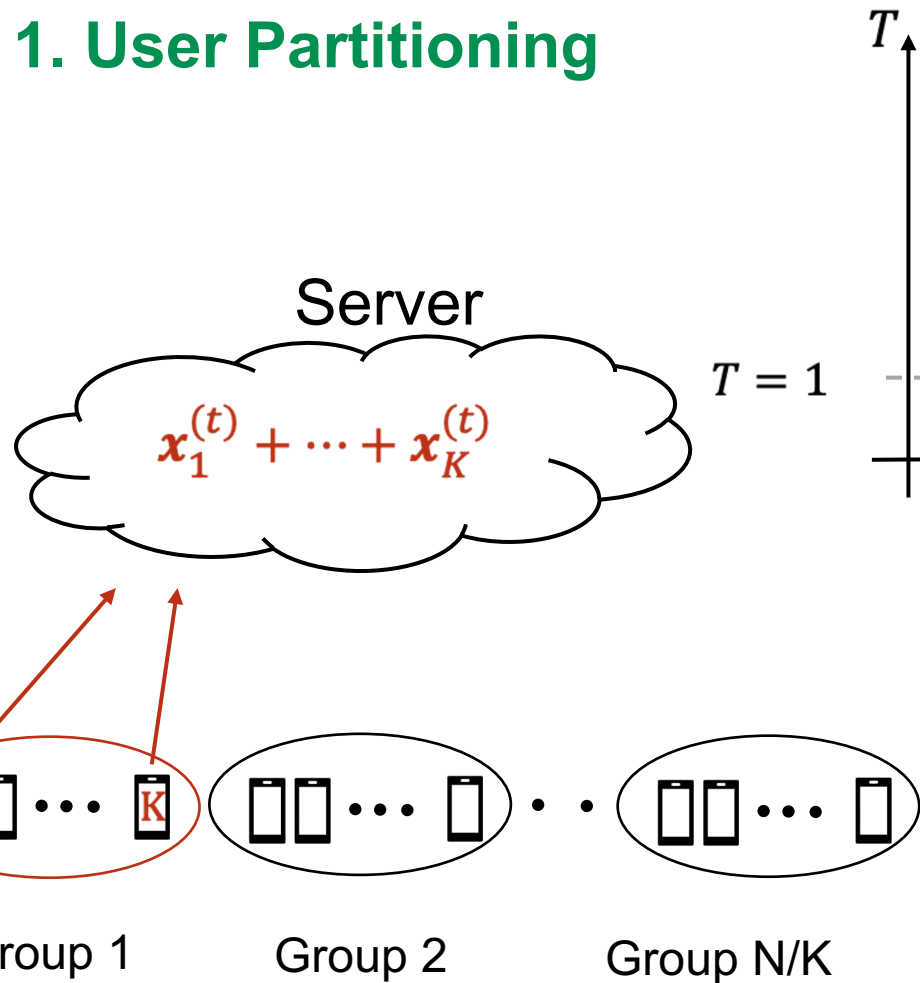
Group 2

Group N/K

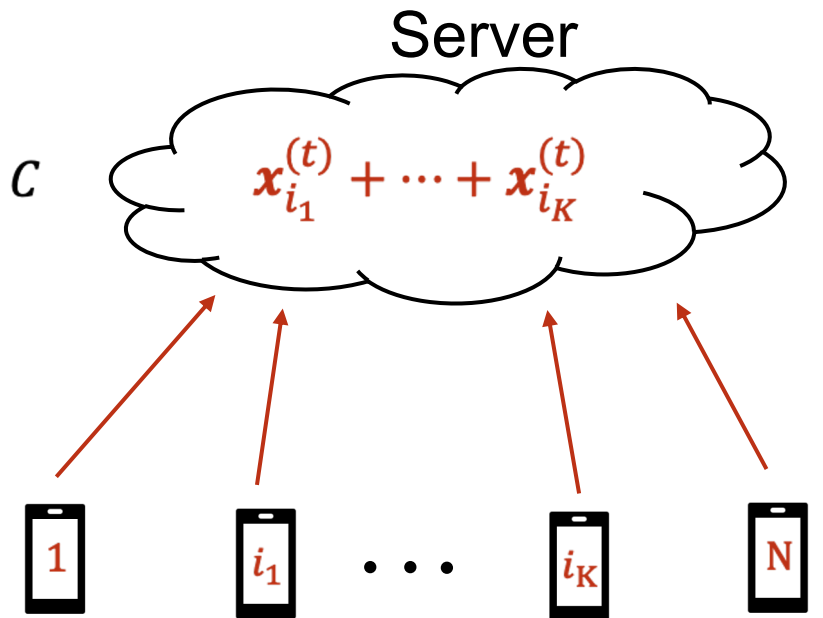
Metric 2: Average Aggregation Cardinality (C)

Aggregation Cardinality C = Average number of participating users over all rounds

1. User Partitioning



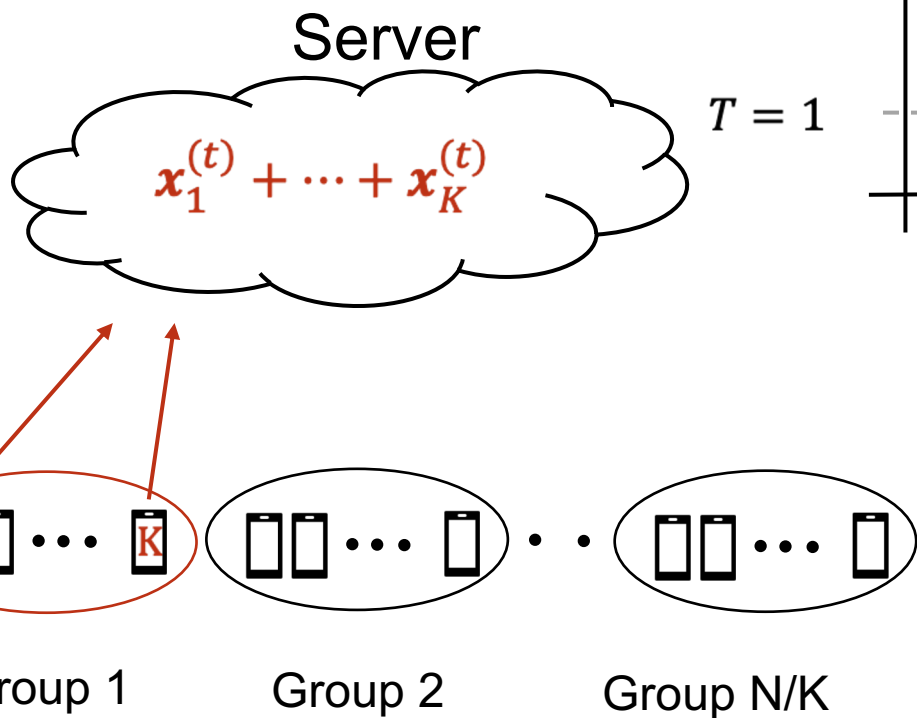
2. Random Selection



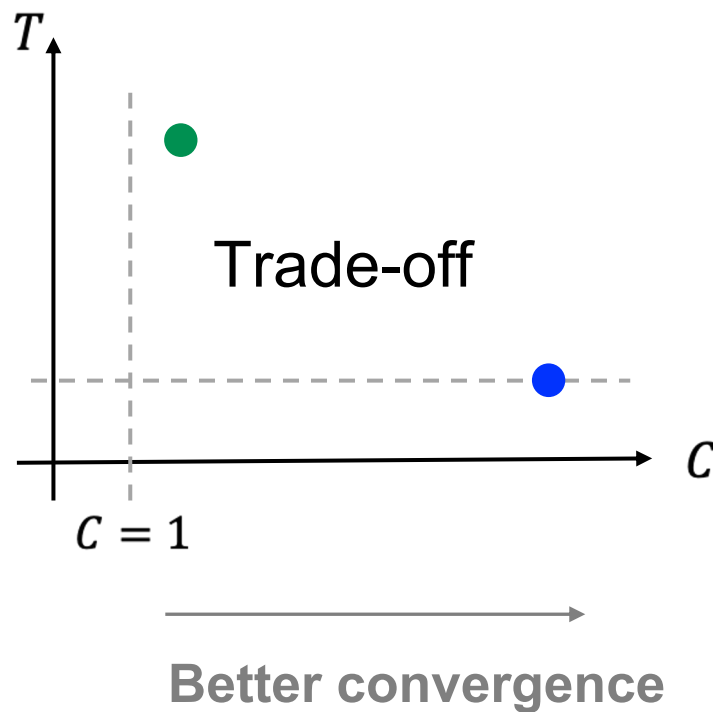

Metric 2: Average Aggregation Cardinality (C)

Aggregation Cardinality C = Average number of participating users over all rounds

1. User Partitioning




2. Random Selection



Metric 3: Aggregation Fairness Gap (F)

- Aggregation Fairness Gap F

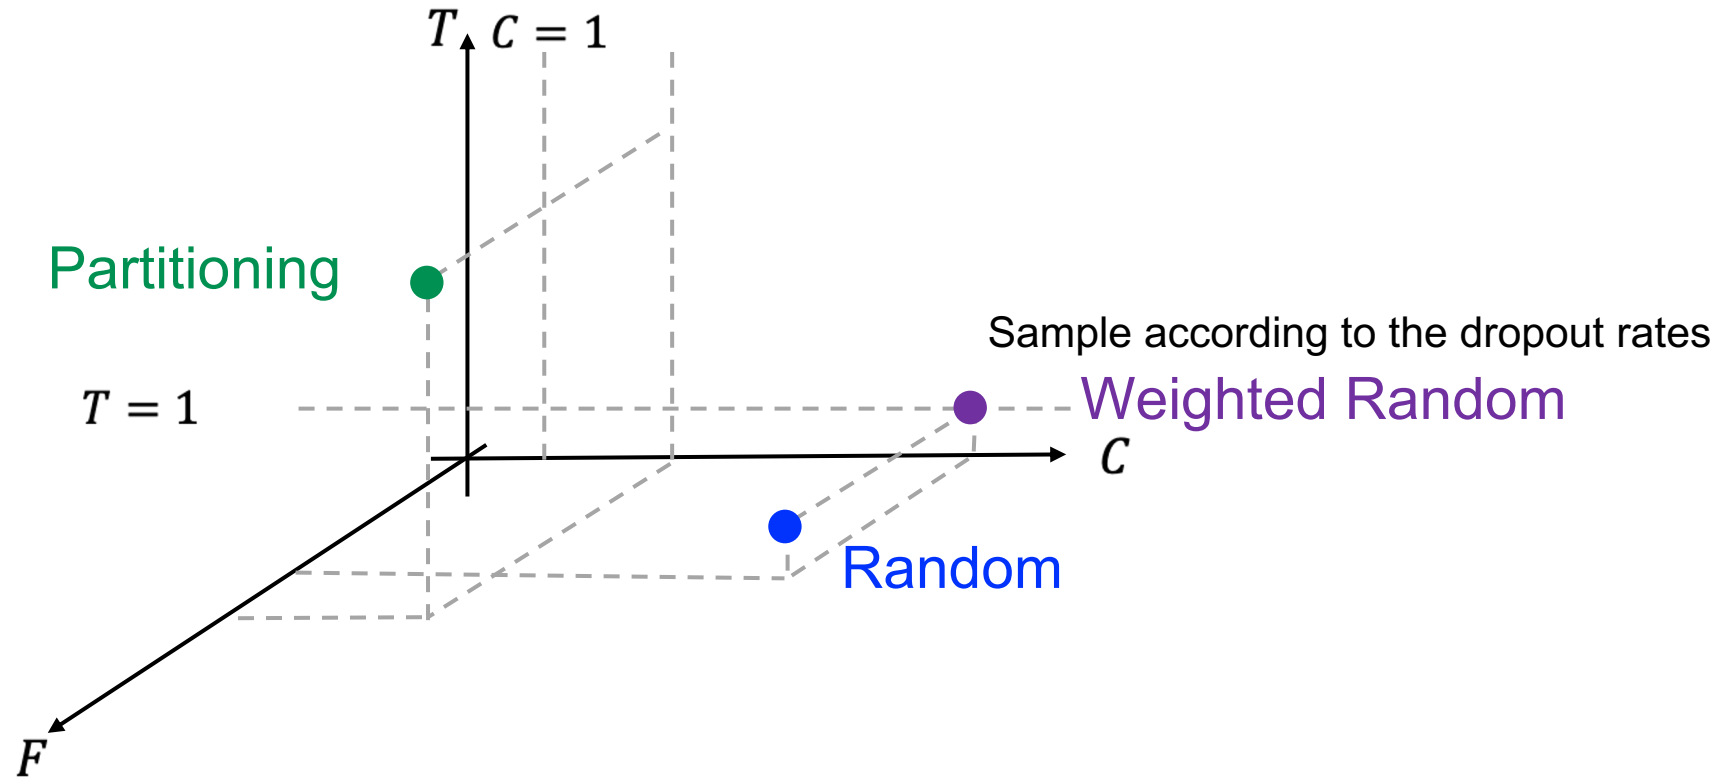
$F = \text{max. average participation frequency} - \text{min. average participation frequency}$



Metric 3: Aggregation Fairness Gap (F)

- Aggregation Fairness Gap F

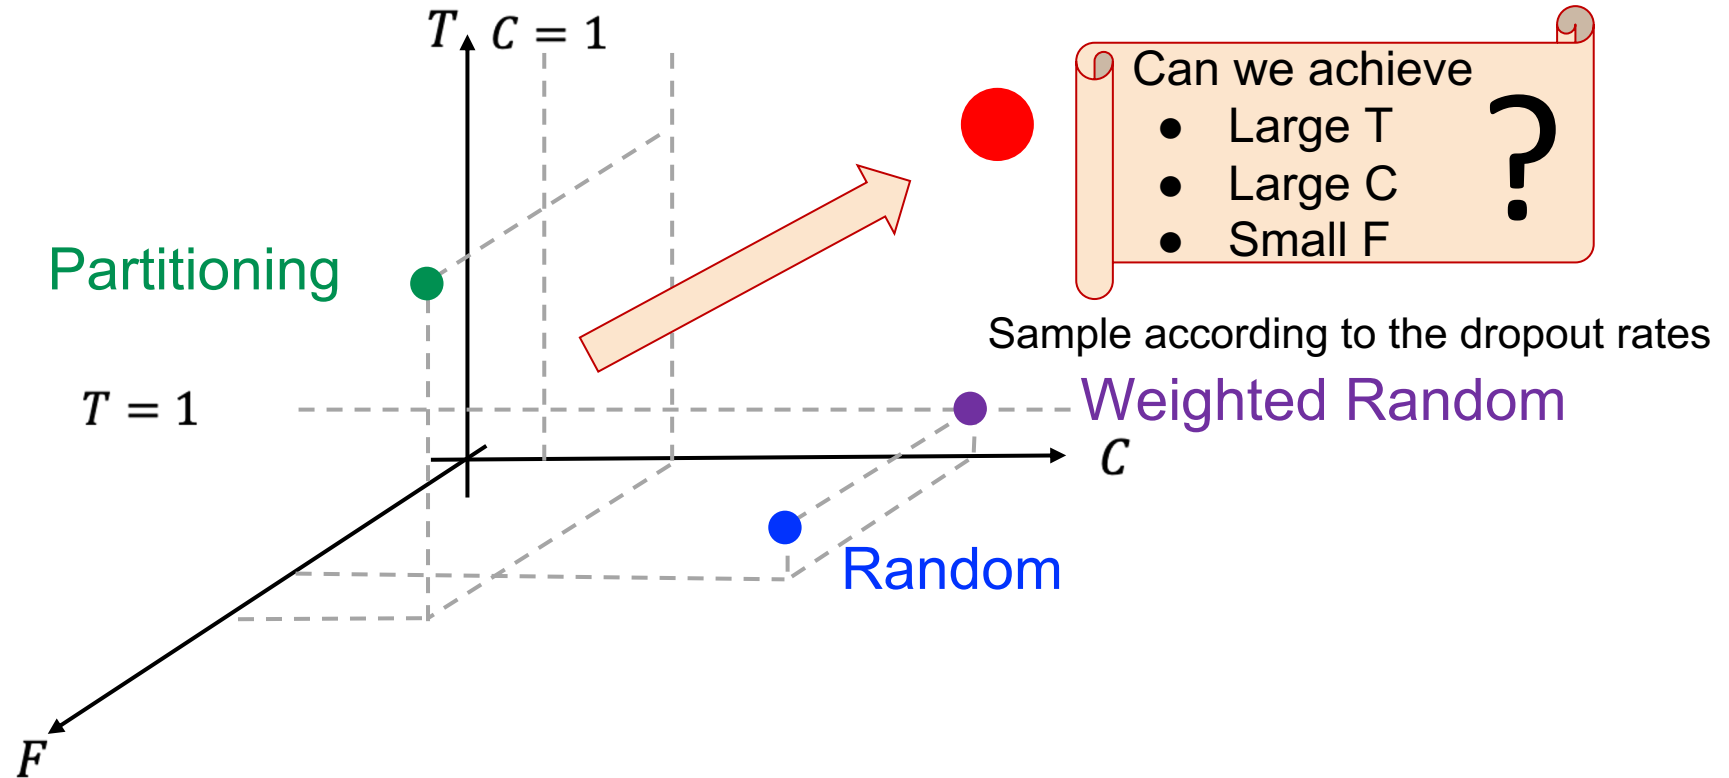
$F = \text{max. average participation frequency} - \text{min. average participation frequency}$



Metric 3: Aggregation Fairness Gap (F)

- Aggregation Fairness Gap F

$F = \text{max. average participation frequency} - \text{min. average participation frequency}$



Proposed Approach: Multi-RoundSecAgg

1) Batch Partitioning

- Input: N , $K \leq N$, $1 \leq T \leq K$
- Output: A family of K -user sets satisfying the multi-round privacy T .
This family is represented by a matrix \mathbf{B} .

2) Available batch selection to guarantee fairness

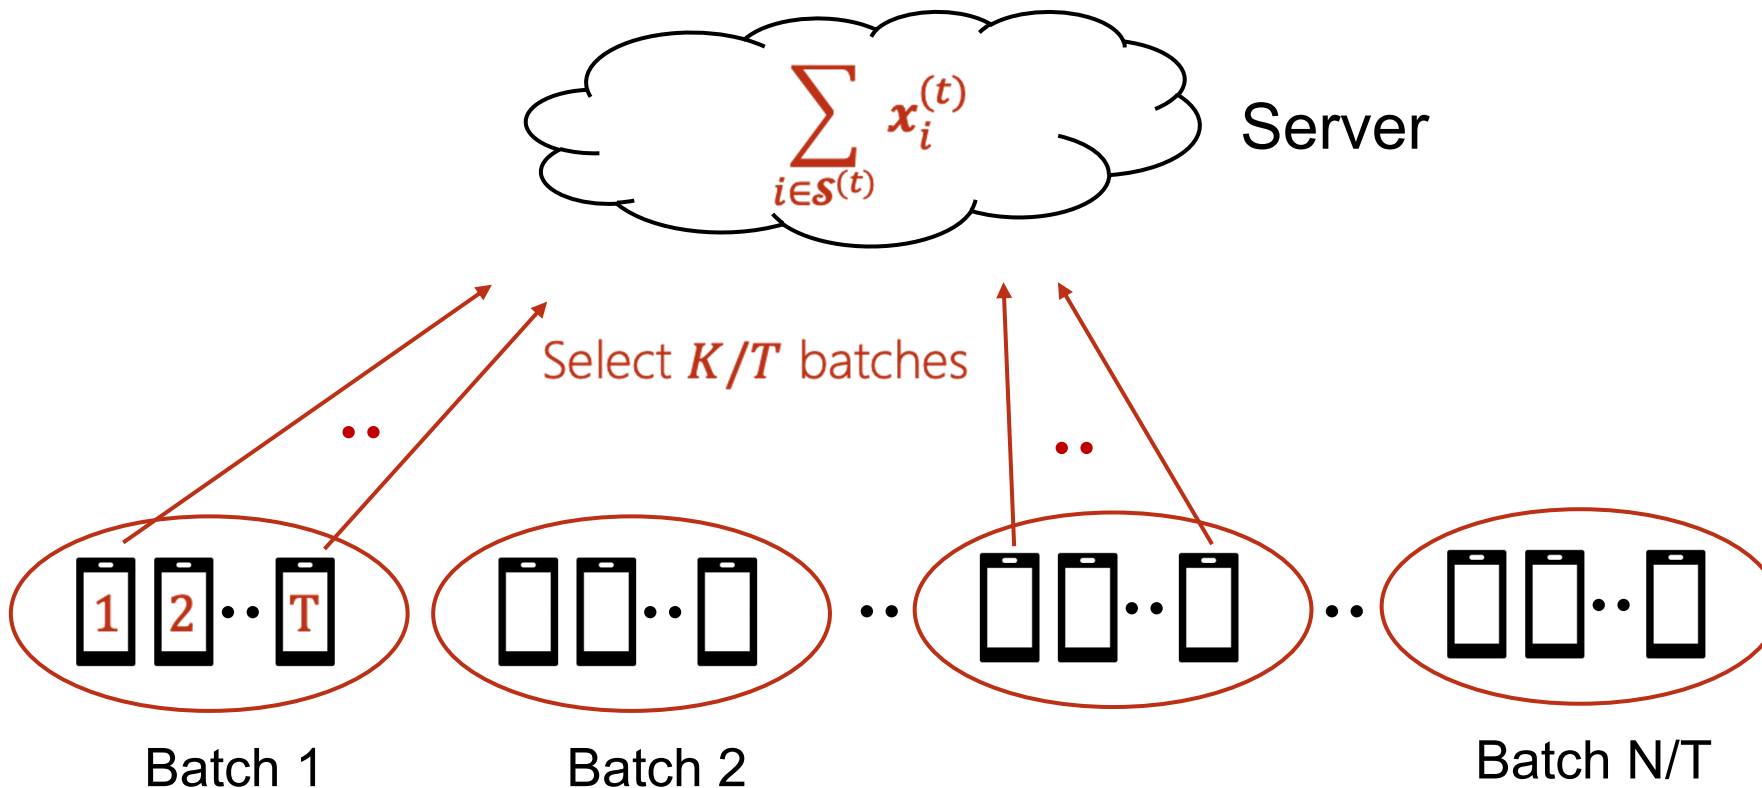
- Input: Set of available users at round t and \mathbf{B} .
- Output: Set of users that will participate at round t .

Multi-RoundSecAgg

1) Batch Partitioning

Idea: Partition users into T -user batches; allow selection of any K/T available batches

This results in a family of $\mathbf{R}_{BP} = \binom{N/T}{K/T}$ sets.



Multi-RoundSecAgg

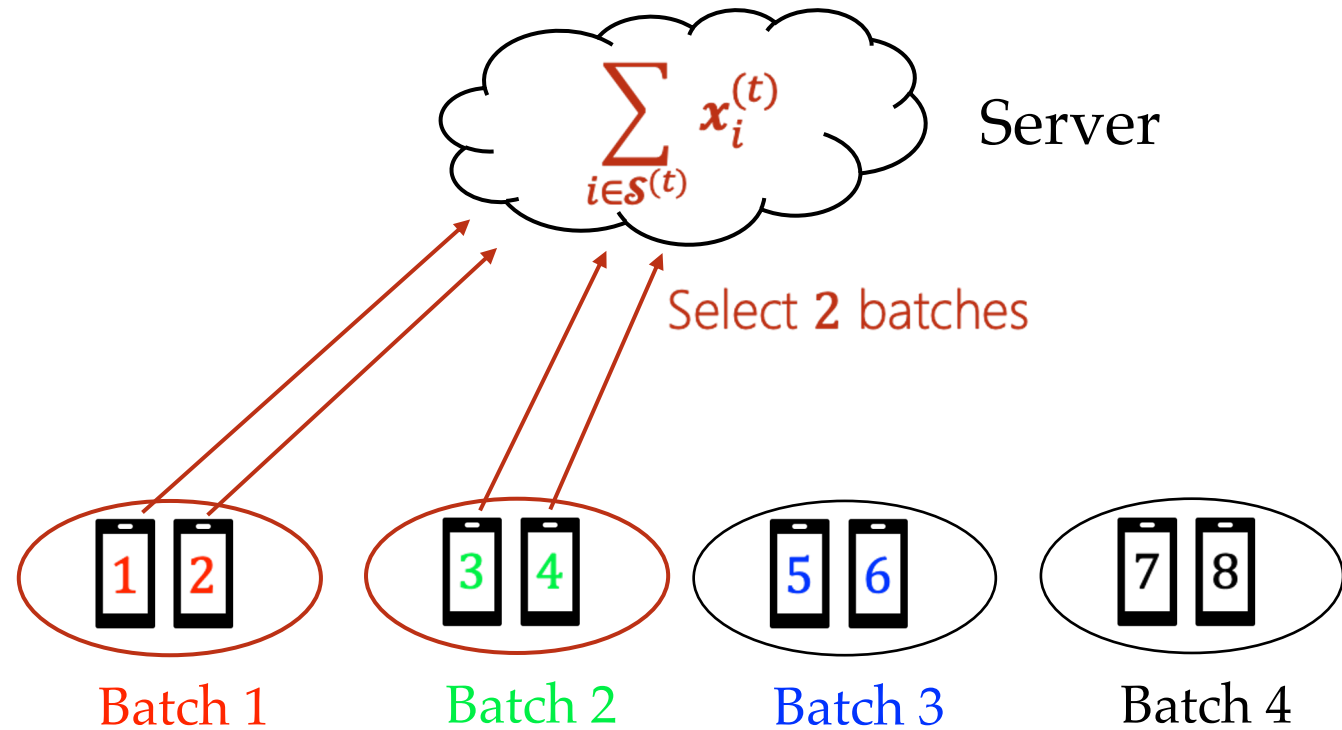
1) Batch Partitioning

Example ($N = 8, K = 4$ & $T = 2$)

$$R_{BP} = \binom{N/T}{K/T} = 6 \text{ sets}$$

$$B = \begin{pmatrix} 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \end{pmatrix}$$

Batch 1 Batch 2 Batch 3 Batch 4



Multi-RoundSecAgg


1) Batch Partitioning

Example ($N = 8, K = 4$ & $T = 2$)

$$R_{BP} = \binom{N/T}{K/T} = 6 \text{ sets}$$

$$B = \begin{pmatrix} 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \end{pmatrix}$$

Batch 1 Batch 2 Batch 3 Batch 4

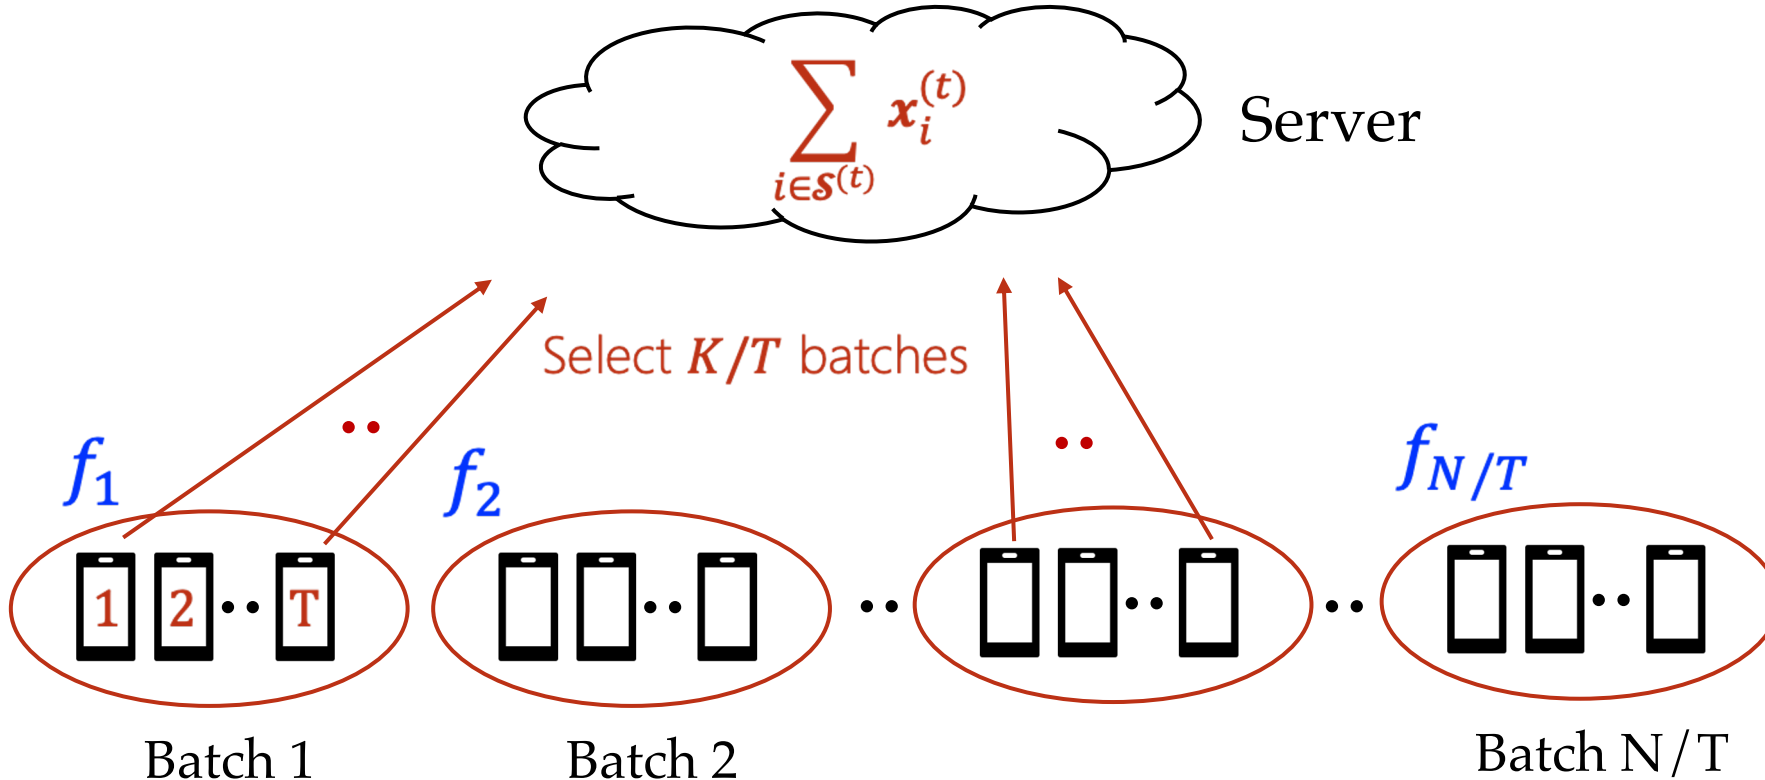


Theorem: aggregated models in the same batch can't be separated across different rounds even through non-linear mixtures of received aggregates.

Multi-RoundSecAgg

2) Available batch selection to guarantee fairness

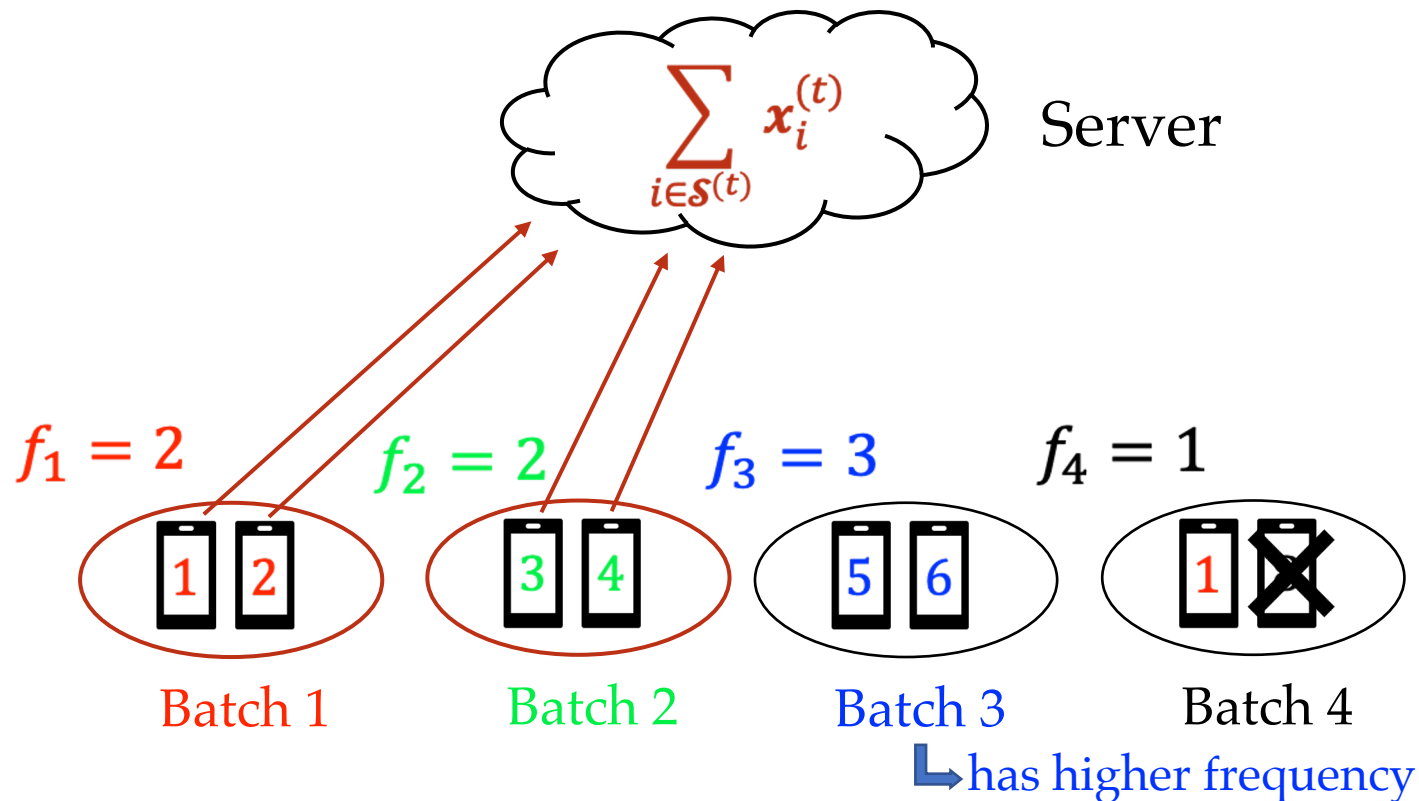
- Input: Set of available users at round t and \mathbf{B} .
- Output: Set of users that will participate at round t .
- Idea: Select based on the minimum frequency of participation.



Multi-RoundSecAgg




2) Available batch selection to guarantee fairness

- Input: Set of available users at round t and \mathbf{B} .
- Output: Set of users that will participate at round t .
- Idea: Select based on the minimum frequency of participation.



An Illustrative Experiment

- Experiment (N=40 users)
 - MNIST Dataset & Non-IID Setting.
 - K=8 users are selected at random at each round.
 - Dropout probability $p_i \sim \{0.1, 0.2, 0.3, 0.4, 0.5\}$
 - The server estimates the individual gradients through least squares.



- **Random Selection**
Reconstruction Error < 0.005 for many users

- **Multi-RoundSecAgg (T=2)**
Reconstruction Error > 0.25 for all users

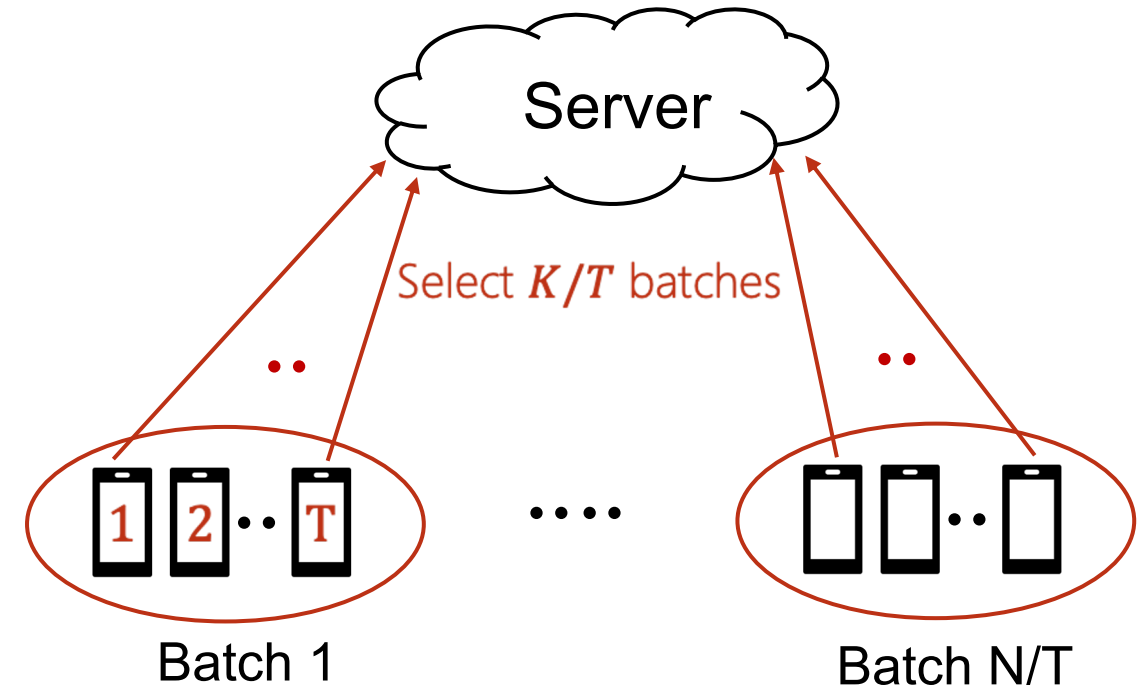
Multi-RoundSecAgg Theoretical Guarantees

Theorem: Multi-RoundSecAgg with parameters $N, K \leq N$ and $1 \leq T \leq K$ ensures

1. a multi-round privacy $1 \leq T \leq K$,
2. an aggregation fairness gap $F = 0$, and
3. an average aggregation cardinality C

$$C(T) = K \left(1 - \sum_{i=N/T-K/T+1}^{N/T} \binom{N/T}{i} q^i (1-q)^{N/T-i} \right),$$

$q = 1 - (1-p)^T$, p : dropout probability



Each batch is not available with probability q

Multi-RoundSecAgg Theoretical Guarantees


Theorem: Multi-RoundSecAgg with parameters $N, K \leq N$ and $1 \leq T \leq K$ ensures

1. a multi-round privacy $1 \leq T \leq K$,
2. an aggregation fairness gap $F = 0$, and
3. an average aggregation cardinality C

$$C(T) = K \left(1 - \sum_{i=N/T-K/T+1}^{N/T} \binom{N/T}{i} q^i (1-q)^{N/T-i} \right),$$

$$q = 1 - (1-p)^T, p: \text{dropout probability}$$

Example ($N = 120, K = N/10 = 12, p = 0.2$)
for $T = N/20 = 6$, we have $C = 11.77 \approx N/10$




Each batch is not available with probability q

Multi-RoundSecAgg Theoretical Guarantees

Theorem: Multi-RoundSecAgg with parameters $N, K \leq N$ and $1 \leq T \leq K$ ensures

1. a multi-round privacy $1 \leq T \leq K$,
2. an aggregation fairness gap $F = 0$, and
3. an average aggregation cardinality C

Better
convergence



$N = 120$ and $K = 12$

Trade-off between “Multi-round Privacy Guarantee” & “Average Aggregation Cardinality”

Multi-RoundSecAgg Convergence Guarantees

Assumptions

1. The loss functions L_1, L_2, \dots, L_N are ρ -smooth.
2. The loss functions L_1, L_2, \dots, L_N are μ -strongly convex.
3. The variance of the stochastic gradients at user i is bounded by σ_i^2 .
4. The expected squared norm of the stochastic gradients is uniformly bounded by G^2 .

$$E[L(\mathbf{x}^{(J)})] - L^* \leq \frac{\rho}{\gamma + \frac{C}{K}JE - 1} \left(\frac{2(\alpha + \beta)}{\mu^2} + \frac{\gamma}{2} E[\|\mathbf{x}^{(0)} - \mathbf{x}^*\|] \right) ,$$


C controls the convergence rate

Experiments


Setup

- $N = 120$ users, $K = 12$ users.
- Dataset: CIFAR-10
- Architecture: LeNet
- Dropout probability, $p_i \sim \{0.1, 0.2, 0.3, 0.4, 0.5\}$

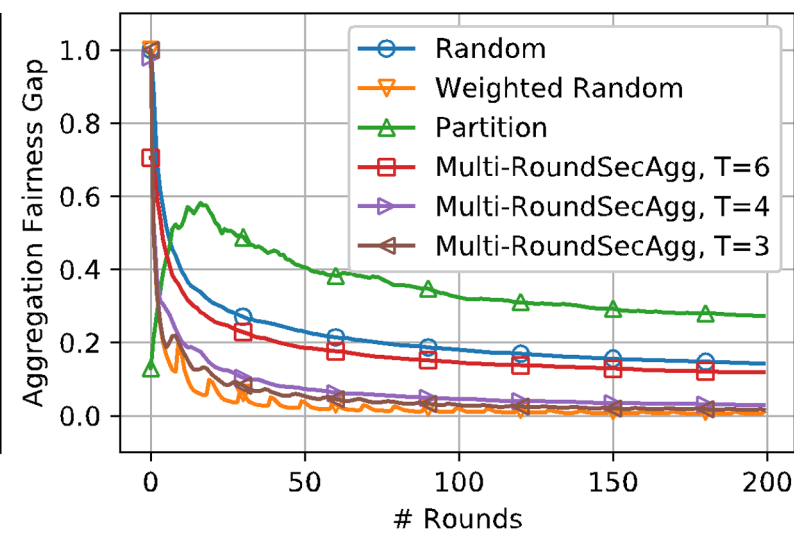
Key Metrics



Multi-round privacy guarantee (T)



Average aggregation cardinality (C)




Aggregation fairness gap (F)


Experiments

Setup


- $N = 120$ users, $K = 12$ users.
- Dataset: CIFAR-10
- Architecture: LeNet
- Dropout probability, $p_i \sim \{0.1, 0.2, 0.3, 0.4, 0.5\}$



(a) I.I.D Data Distribution



(b) Non I.I.D Data Distribution



(c) Trade-off between multi-round privacy guarantee & average aggregation cardinality

Concluding Remarks

- Random user selection in FL can lead to serious privacy leakage
- MultiRoundSecAgg is the first scheme for mitigating this challenge
- MultiRoundSecAgg reveals an interesting tradeoff between “privacy” and “convergence rate” in FL

Concluding Remarks

- Random user selection in FL can lead to serious privacy leakage
- MultiRoundSecAgg is the first scheme for mitigating this challenge
- MultiRoundSecAgg reveals an interesting tradeoff between “privacy” and “convergence rate” in FL
- Potential future directions
 - Our metric for multi-round privacy is very strong. Careful relaxations may lead to substantial improvements (e.g., in aggregation cardinality)
 - Formalizing a fundamental trade-off between privacy and convergence-rate in FL?
 - Our protocol guarantees that only an aggregate of models can be learned. How to bound privacy leakage from aggregate models?

$$\text{🔒 } x_1^{(t)} + x_2^{(t)} + \dots + x_N^{(t)}$$



Questions?
Thank you

Additional Slides

Optimality of Multi-RoundSecAgg

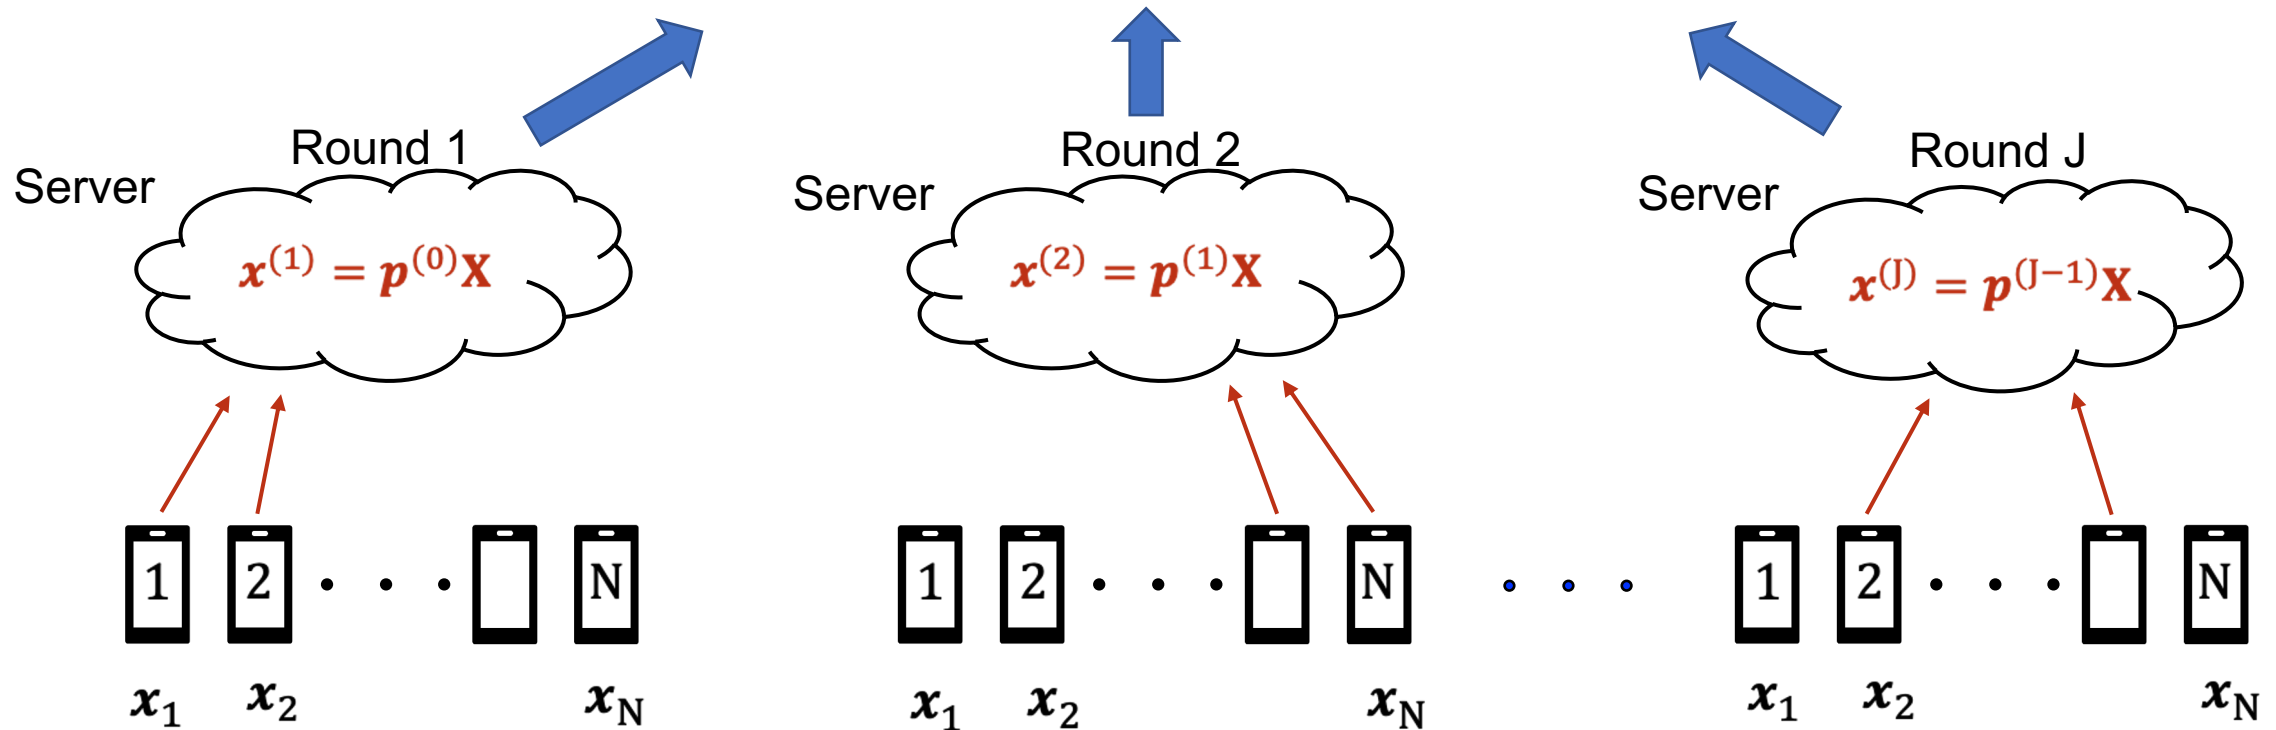
- Any strategy satisfying the multi-round privacy guarantee must have a batch partitioning structure.
- For given $N, K \leq N/2$ and T , any strategy satisfying a multi-round privacy \mathbf{T} can have at most R_{\max} user sets

$$R_{\max} \leq \binom{N/T}{K/T} = R_{\text{BP}} \quad \text{number of sets in BP}$$

Our Multi-round Privacy is Strong

- A multi-round privacy T requires that any non-zero partial sum that the server can reconstruct to be of the form

$$a_1 \sum_{j \in \mathcal{S}_1} \mathbf{x}_j + a_2 \sum_{j \in \mathcal{S}_2} \mathbf{x}_j + \dots + a_n \sum_{j \in \mathcal{S}_n} \mathbf{x}_j, \text{ where } |\mathcal{S}_i| \geq T.$$

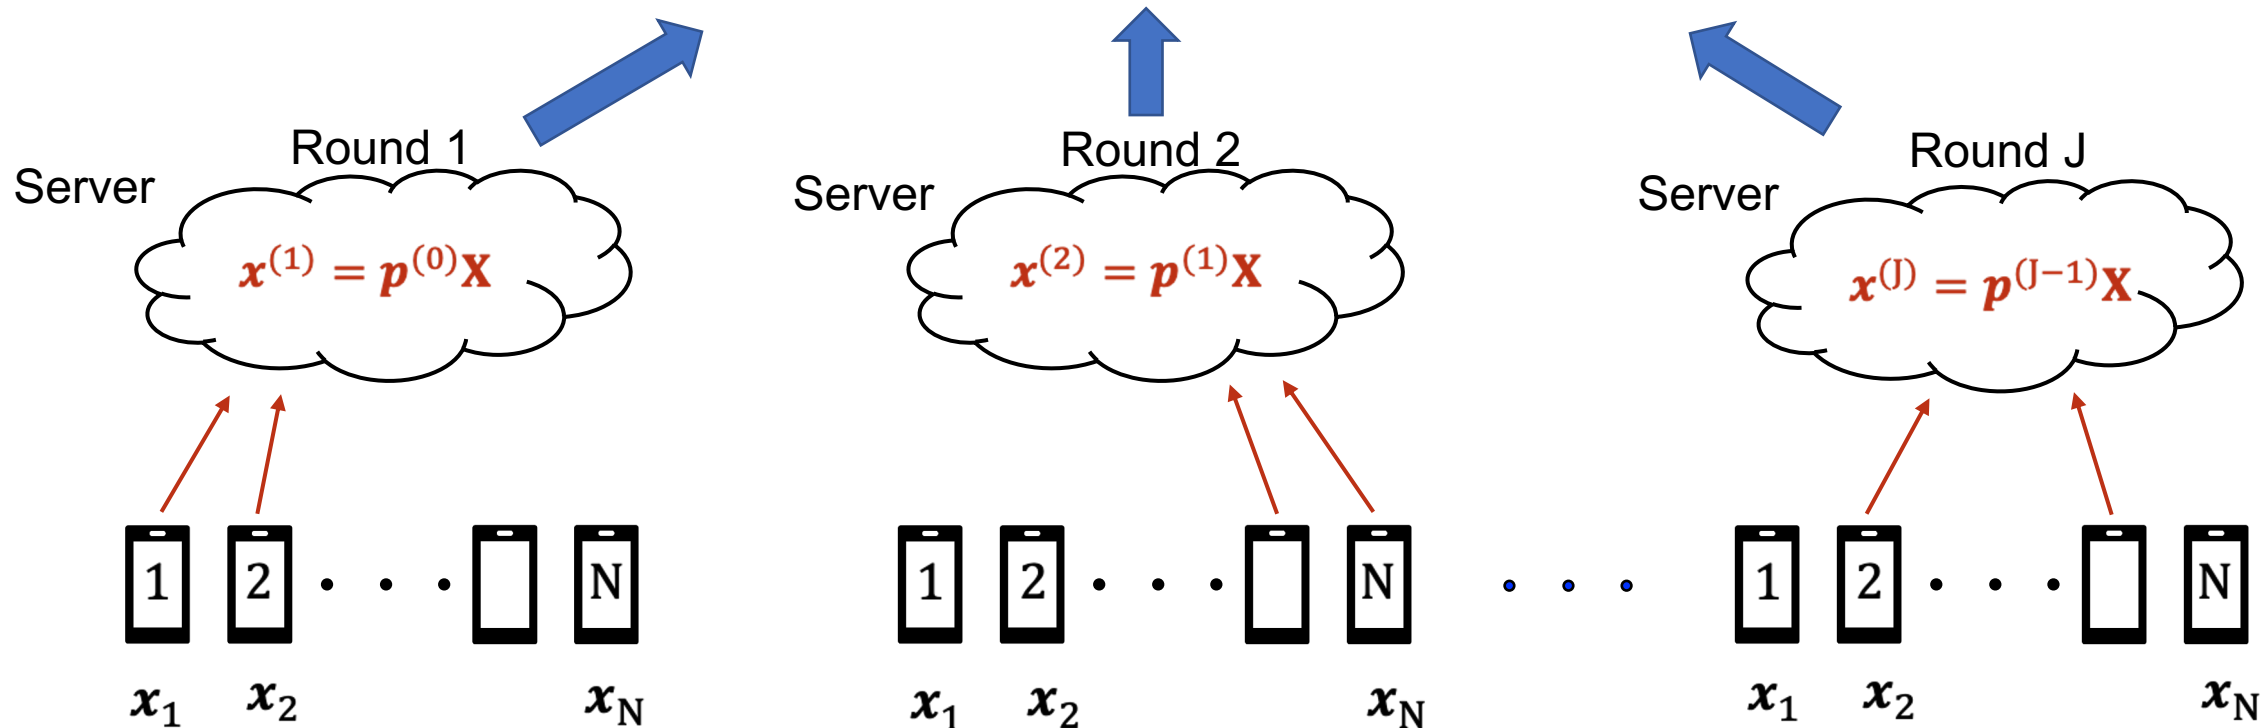


Each group has the same coefficient.

Relaxed Multi-round Privacy

- The relaxed multi-round privacy T requires that any non-zero partial sum that the server can reconstruct to be of the form

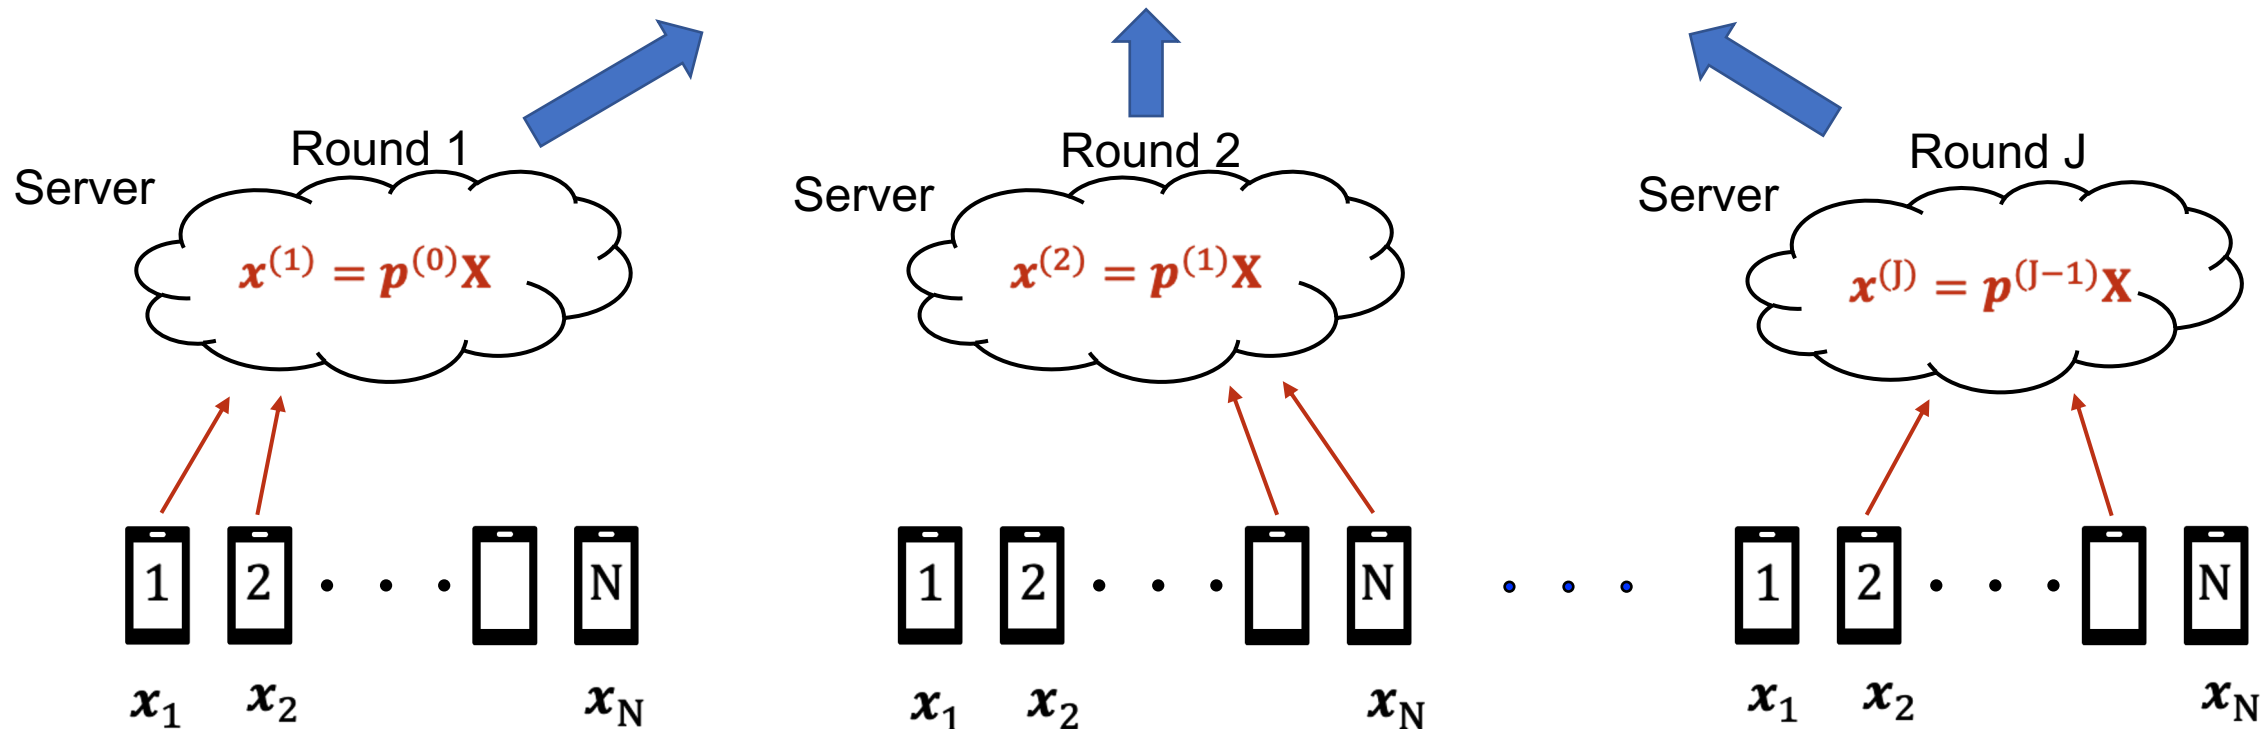
$$\sum_{j \in \mathcal{S}} a_j \mathbf{x}_j, \text{ where } |\mathcal{S}| \geq T$$



Relaxed Multi-round Privacy

- The relaxed multi-round privacy T requires that any non-zero partial sum that the server can reconstruct to be of the form

$$\sum_{j \in \mathcal{S}} a_j \mathbf{x}_j, \text{ where } |\mathcal{S}| \geq T$$

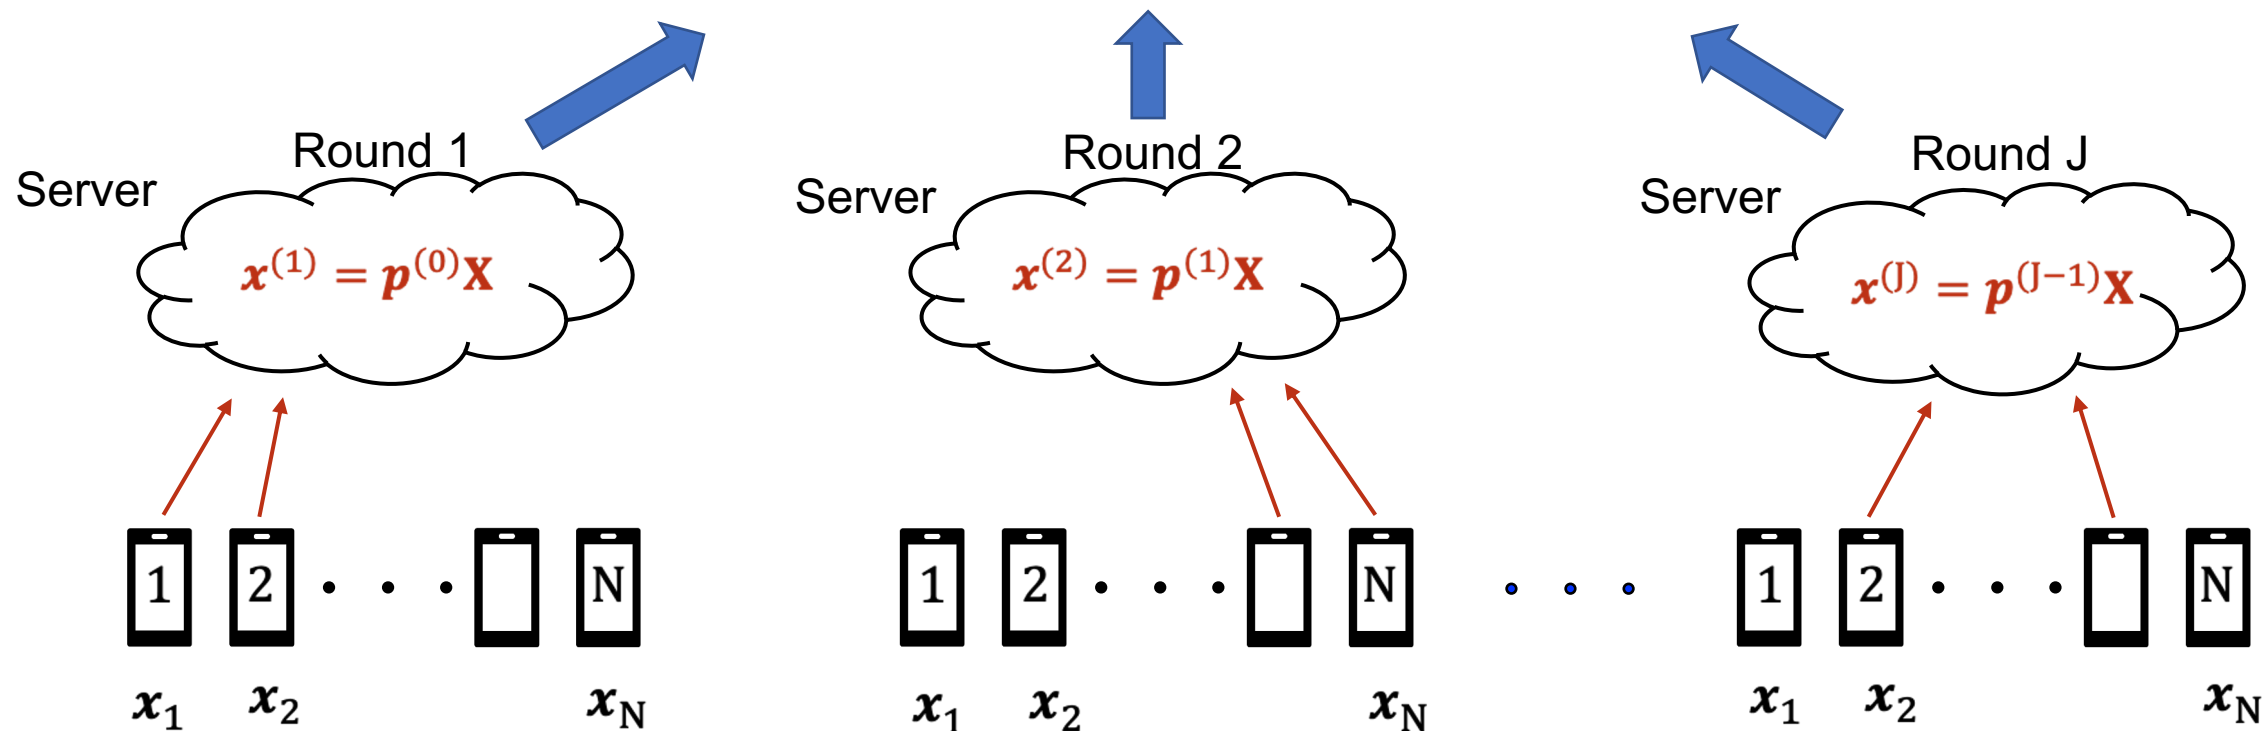


When $T = 2$, this allows reconstructing $a_i \mathbf{x}_i + a_j \mathbf{x}_j$ which can reveal \mathbf{x}_i if $a_i \gg a_j$.

Relaxed Multi-round Privacy

- The relaxed multi-round privacy T requires that any non-zero partial sum that the server can reconstruct to be of the form

$$\sum_{j \in \mathcal{S}} a_j \mathbf{x}_j, \text{ where } |\mathcal{S}| \geq T$$




Batch partitioning is not necessary

Relaxed Multi-round Privacy

- The relaxed multi-round privacy T requires that any non-zero partial sum that the server can reconstruct to be of the form

$$\sum_{j \in \mathcal{S}} a_j \mathbf{x}_j, \text{ where } |\mathcal{S}| \geq T$$



What's the optimal strategy?

Relaxed Multi-round Privacy

- The relaxed multi-round privacy T requires that any non-zero partial sum that the server can reconstruct to be of the form

$$\sum_{j \in \mathcal{S}} a_j \mathbf{x}_j, \text{ where } |\mathcal{S}| \geq T$$

Example (N=6, K=4, T=2)

$$\mathbf{B}_{old} = \begin{pmatrix} 1 & 1 & 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 & 1 \end{pmatrix} \xrightarrow{2 \times} \mathbf{B}_{new} = \begin{pmatrix} 1 & 1 & 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 & 1 \\ 1 & 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 & 1 \\ 0 & 1 & 1 & 1 & 1 & 0 \end{pmatrix}$$

Aggregation Fairness Gap F & Average Aggregation Cardinality C

Aggregation Fairness Gap F

$$F = \max_{i \in [N]} \limsup_{J \rightarrow \infty} \frac{1}{J} \mathbb{E} \left[\sum_{t=0}^{J-1} \{\mathbf{p}^{(t)}\}_i \right] - \min_{i \in [N]} \liminf_{J \rightarrow \infty} \frac{1}{J} \mathbb{E} \left[\sum_{t=0}^{J-1} \{\mathbf{p}^{(t)}\}_i \right]$$

maximum average participation frequency

minimum average participation frequency

Average Aggregation Cardinality C (Average number of participating users)

$$C = \liminf_{J \rightarrow \infty} \frac{1}{J} \mathbb{E} \left[\sum_{t=0}^{J-1} \|\mathbf{p}^{(t)}\|_0 \right]$$

Multi-RoundSecAgg Convergence Guarantees

Assumptions

1. The loss functions L_1, L_2, \dots, L_N are ρ -smooth.
2. The loss functions L_1, L_2, \dots, L_N are μ -strongly convex.
3. The variance of the stochastic gradients at user i is bounded by σ_i^2 .
4. The expected squared norm of the stochastic gradients is uniformly bounded by G^2 .

$$E[L(\mathbf{x}^{(J)})] - L^* \leq \frac{\rho}{\gamma + \frac{C}{K} J E - 1} \left(\frac{2(\alpha + \beta)}{\mu^2} + \frac{\gamma}{2} E[\|\mathbf{x}^{(0)} - \mathbf{x}^*\|] \right),$$

where E is the number of local SGD steps, $\alpha = \frac{1}{N} \sum_{i=1}^N \sigma_i^2 + 6\rho\Gamma + 8(E-1)^2 G^2$,
 $\beta = \frac{4(N-K)E^2 G^2}{K(N-1)}$, $\Gamma = L^* - \sum_{i=1}^N L_i$ and $\gamma = \max\{\frac{8\rho}{\mu}, E\}$

Necessity of Batch Partitioning (BP)

- Any strategy satisfying the multi-round privacy guarantee must have a batch partitioning structure.

Necessity of Batch Partitioning (BP)

- Any strategy satisfying the multi-round privacy guarantee must have a batch partitioning structure.

Proof Idea

Consider any scheme with a matrix $\mathbf{V}_{R \times N}$ and denote the linear combination used by the server by $\mathbf{z}_{1 \times R}$, $\mathbf{z}_i \sim U[0, 1], i \in [R]$.

For the scheme to have privacy T , at least T elements of \mathbf{zV} have to be equal.

Necessity of Batch Partitioning (BP)

- Any strategy satisfying the multi-round privacy guarantee must have a batch partitioning structure.

Proof Idea

Consider any scheme with a matrix $\mathbf{V}_{R \times N}$ and denote the linear combination used by the server by $\mathbf{z}_{1 \times R}$, $\mathbf{z}_i \sim U[0, 1], i \in [R]$.

For the scheme to have privacy T , at least T elements of \mathbf{zV} have to be equal.

This implies that at least T columns of \mathbf{V} are equal as $\mathbf{z}_i \sim U[0, 1]$.

➡ Batch partitioning is necessary

Necessity of Batch Partitioning (BP)

- Any strategy satisfying the multi-round privacy guarantee must have a batch partitioning structure.

Proof Idea

Consider any scheme with a matrix $\mathbf{V}_{R \times N}$ and denote the linear combination used by the server by $\mathbf{z}_{1 \times R}$, $\mathbf{z}_i \sim U[0, 1], i \in [R]$.

For the scheme to have privacy T , at least T elements of \mathbf{zV} have to be equal.

This implies that at least T columns of \mathbf{V} are equal as $\mathbf{z}_i \sim U[0, 1]$.

➡ Batch partitioning is necessary

- For given $N, K \leq N/2$ and T , any strategy satisfying a multi-round privacy T can have at most R_{\max} user sets

$$R_{\max} \leq \binom{N/T}{K/T} = R_{\text{BP}} \text{ number of sets in BP}$$

References

- [1] H Brendan McMahan et al.. **“Communication-efficient learning of deep networks from decentralized data.”**
In Int. Conf. on Artificial Int. and Stat. (AISTATS), pages 1273–1282, 2017.
- [2] Balázs Pejó, and Gergely Biczók. **"Quality Inference in Federated Learning with Secure Aggregation."**
arXiv:2007.06236 (2020).
- [3] M. Lam et al. **"Gradient Disaggregation: Breaking Privacy in Federated Learning by Reconstructing the User Participant Matrix."** arXiv:2106.06089 (2021).