Incentivized Communication for Federated Bandits

Zhepei Wei†, Chuanhao Li†, Haifeng Xu‡, Hongning Wang†
†University of Virginia, ‡University of Chicago

Motivation

Typical Federated Bandits

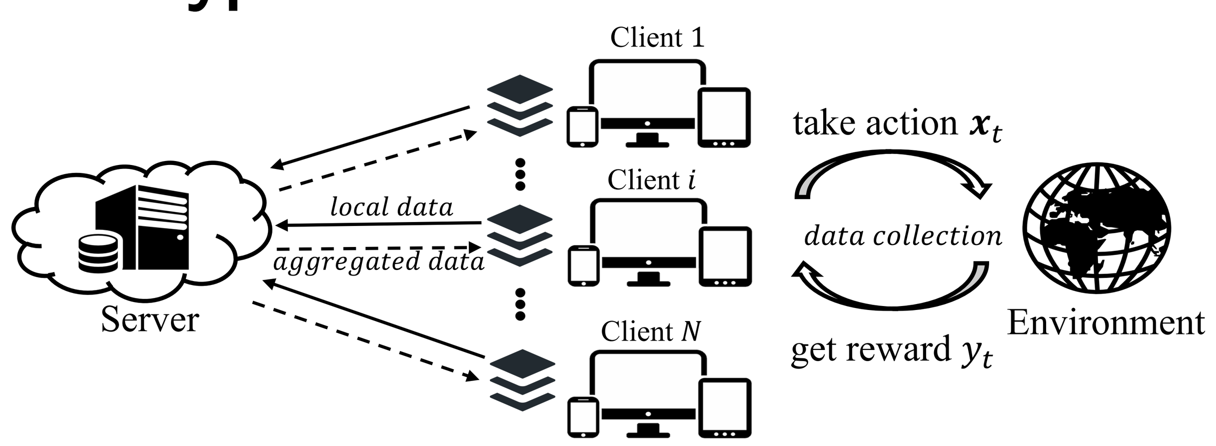

For time $t = 1, 2, \dots, T$
For client $i = 1, 2, \dots, N$
- Client i takes **action** x_t from action set $\mathcal{A}_t$ and observes **reward** $y_t = f(x_t) + \eta_t$
- Communication between server and clients

Focus: efficient communication protocol design that trades off communication cost and regret.

$$R_T = \sum_{t=1}^{NT} r_t, \text{ where } r_t = \max_{x\in\mathcal{A}_t} f(x) - f(x_t)$$

Unveiling the Achilles' Heel: existing protocols essentially require/assume full client engagement whenever communication is triggered, **however, what if clients are reluctant to share data and opt-out?**

Problem Formulation

Incentivized Federated Bandits

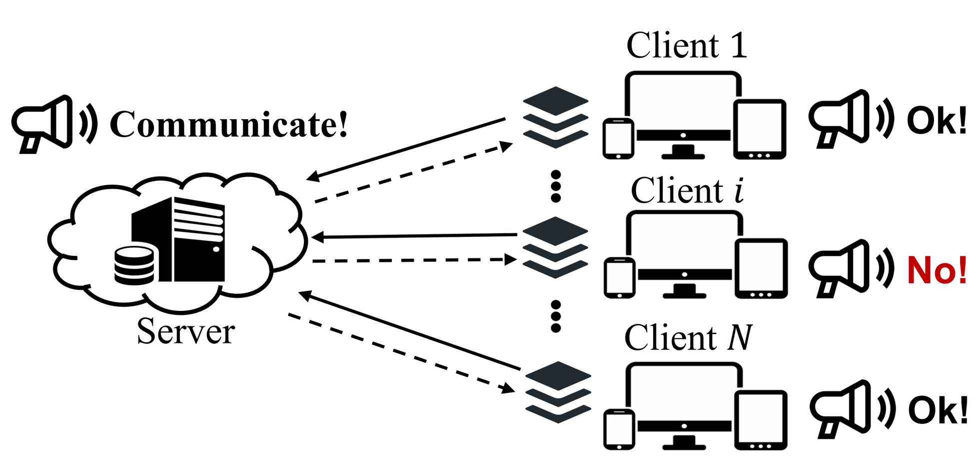

Incentivized Problem Setting: clients are self-interested, and will not share their data with the server unless the benefits outweigh any potential loss of sharing, e.g., privacy breaches. This is characterized by:

- Client decides whether to share data
- Server can motivate clients by providing incentives

Research Question: how to design an incentivized communication protocol that balances multiple objectives, i.e., achieving nearly-optimal regret, with reasonable communication and incentive costs?

Payment-Free Design: Data as Incentive

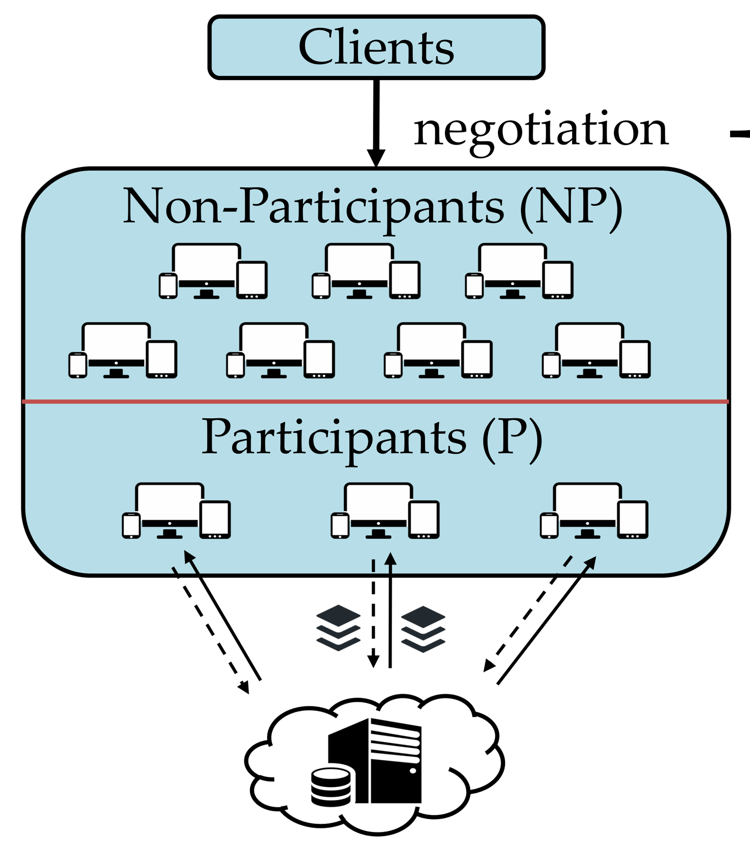

Server: if you share data, I will:
- Offer my reserved data and other participants' uploads

Client: I only care about myself, I will:
- Participate, if your offer exceeds my data sharing cost
- Not participate, otherwise

Data valuation

server's offer for client i — client i's data sharing cost

$$\frac{\det\left(\sum_{j:\{\Delta V_{j,t}\in S\}\wedge\{j\neq i\}} \Delta V_{j,t} + \Delta V_{-i,t} + V_{i,t} + \lambda I\right)}{\det(V_{i,t} + \lambda I)} > D_p \cdot \mathbb{I}(\Delta V_{i,t} \neq \mathbf{0})$$

client i's local data

Regret Bound

$$r_t \le 2\alpha_{i_t,t-1}\sqrt{\mathbf{x}_t^\top \widetilde{V}_{t-1}^{-1}\mathbf{x}_t}\cdot\sqrt{\frac{\det(\widetilde{V}_{t-1})}{\det(V_{i_t,t-1})}} = O\left(\sqrt{d\log\frac{T}{\delta}}\right)\cdot\|\mathbf{x}_t\|_{\widetilde{V}_{t-1}^{-1}}\cdot\sqrt{\frac{\det(\widetilde{V}_{t-1})}{\det(V_{i_t,t-1})}}$$

We proved that, to achieve near-optimal regret, it is required that the shared data at each communication round is at least above a threshold compared to all available data in the system.

However, as this payment-free data exchange cannot force participation, **it can not guarantee regret**.

Regret not Guaranteed :
Regret Guaranteed:

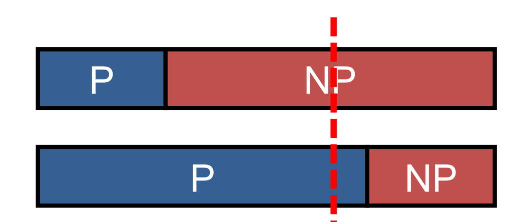

Payment-Efficient Design: Money as Additional Incentive

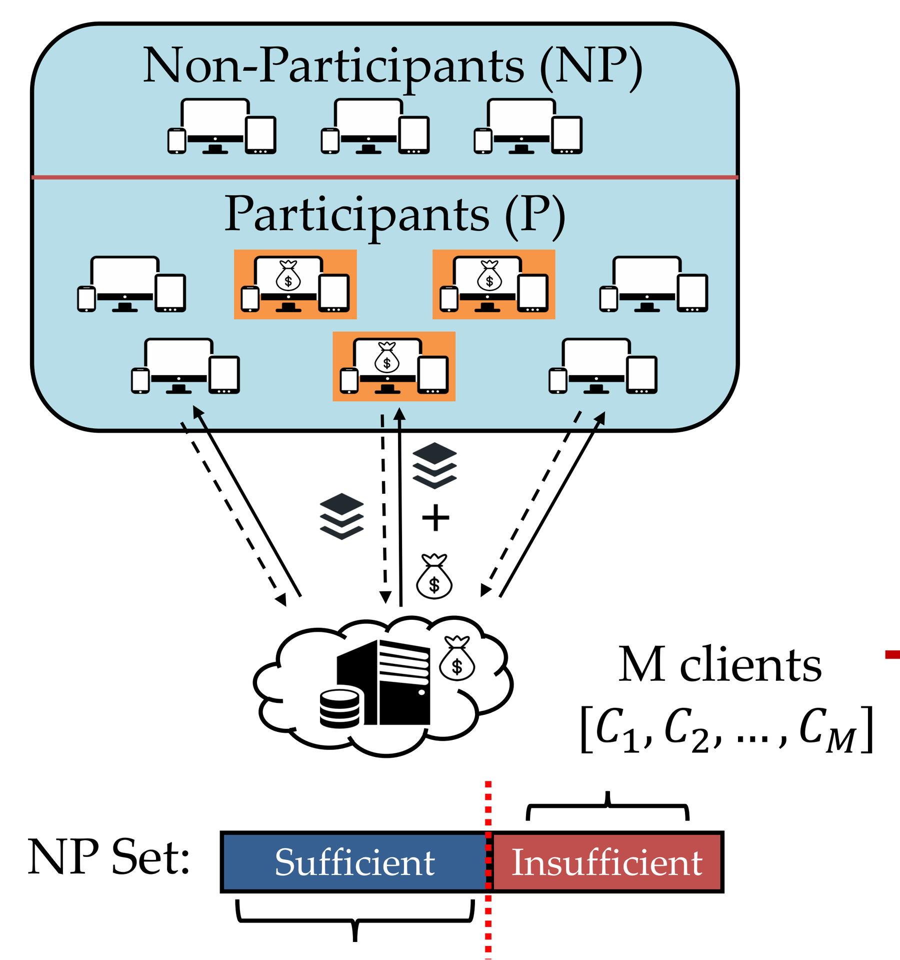

Any client in sufficient list is an immediate valid choice, to save incentive, we can choose the first client in this list as our **last resort**.

Goal: passing the threshold while minimizing total payment.

How to choose clients?

➢ Step 1. Rank clients by their potential contribution, i.e., how much it can help the server to pass the threshold.

➢ Step 2. Check if any combination of $n \le M$ clients within insufficient set can pass the threshold with less incentive cost than our last resort. **Specifically, we always start from the clients with largest contribution.**

- Iteratively choose most contributing n clients, i.e., $[C_{M-n+1}, \dots, C_M]$:
  - If having n (initially, n = 1) clients requires more incentive than last resort, then **terminate**. Otherwise, if resulting participant set pass the threshold, then **terminate**;
  - If having n clients is not enough to pass the threshold, re-rank the contribution list with n clients committed, then **increase $n = n + 1$**.

$O(N)$ Complexity

Theoretical & Empirical Results

We prove that, the proposed payment-efficient solution achieves **near-optimal regret** $R_T = O(d\sqrt{T}\log T)$, with **communication cost** $C_T = O(d^3N^2\log T)$ and **incentive cost** $M_T = O\left(\max D_p \times P \times N - \sum_{i=1}^{N} P_i \times \left(\frac{\det \lambda I}{\det V_T}\right)^{1/P_i}\right)$, where P_i is the number of epochs client i get paid, P is the number of epochs.

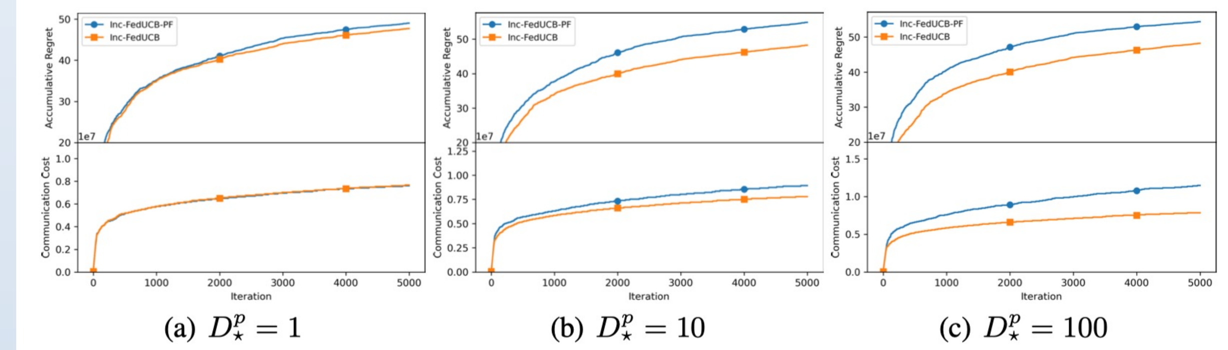

Figure 1: Comparison between payment-free vs. payment-efficient incentive designs.

$d = 25, K = 25$		DisLinUCB	Inc-FedUCB ($\beta = 1$)	Inc-FedUCB ($\beta = 0.7$)	Inc-FedUCB ($\beta = 0.3$)
$T = 5,000, N = 50, D_\star^p = 0$	Regret (Acc.)	48.46	48.46	48.46 ($\Delta = 0\%$)	48.46 ($\Delta = 0\%$)
	Commu. Cost	7,605,000	7,605,000	7,605,000 ($\Delta = 0\%$)	7,605,000 ($\Delta = 0\%$)
	Pay. Cost	\	0	0 ($\Delta = 0\%$)	0 ($\Delta = 0\%$)
$T = 5,000, N = 50, D_\star^p = 1$	Regret (Acc.)	\	48.46	47.70 ($\Delta - 1.6\%$)	48.38 ($\Delta - 0.2\%$)
	Commu. Cost	\	7,605,000	7,668,825 ($\Delta + 0.8\%$)	7,733,575 ($\Delta + 1.7\%$)
	Pay. Cost	\	75.12	60.94 ($\Delta - 18.9\%$)	22.34 ($\Delta - 70.3\%$)
$T = 5,000, N = 50, D_\star^p = 10$	Regret (Acc.)	\	48.46	48.21 ($\Delta - 0.5\%$)	47.55 ($\Delta - 1.9\%$)
	Commu. Cost	\	7,605,000	7,779,425 ($\Delta + 2.3\%$)	8,599,950 ($\Delta + 13\%$)
	Pay. Cost	\	12,819.61	9,050.61 ($\Delta - 29.4\%$)	4,859.17 ($\Delta - 62.1\%$)
$T = 5,000, N = 50, D_\star^p = 100$	Regret (Acc.)	\	48.46	48.22 ($\Delta - 0.5\%$)	48.44 ($\Delta - 0.1\%$)
	Commu. Cost	\	7,605,000	7,842,775 ($\Delta + 3.1\%$)	8,718,425 ($\Delta + 14.6\%$)
	Pay. Cost	\	190,882.45	133,426.01 ($\Delta - 30.1\%$)	88,893.78 ($\Delta - 53.4\%$)

Table 1: Study on Hyper-Parameter of Inc-FedUCB and Environment

Future Work

New Challenge: some adversarial clients may misreport their data sharing costs, and take advantage of the server to increase their utility.

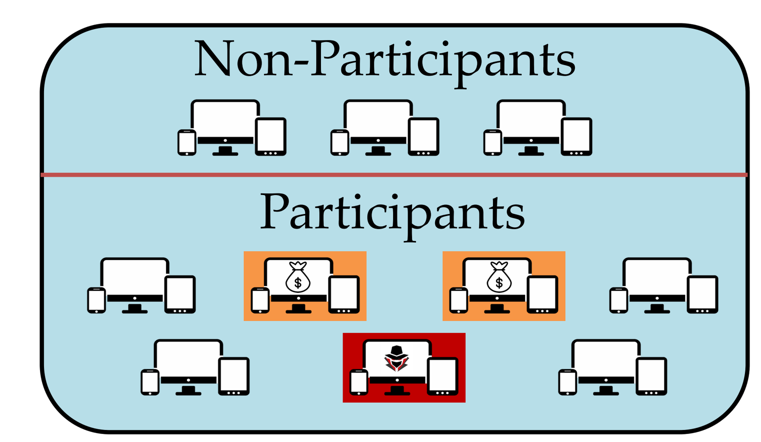

Research Question: how can we incentivize clients in a way that encourages them to truthfully report their costs in their best interest?

Acknowledgement

National Science Foundation

This work is supported by National Science Foundation (NSF) (No. IIS-2213700).