Federated Bandit: A Gossiping Approach

Zhaowei Zhu[∗†], Jingxuan Zhu[∗§], Ji Liu[§], Yang Liu[†]

[†]UC Santa Cruz, [§]Stony Brook University

{zwzhu,yangliu}@ucsc.edu,{jingxuan.zhu,ji.liu}@stonybrook.edu

(*Equal contributions)

UNIVERSITY OF CALIFORNIA SANTA CRUZ

Stony Brook University

Motivating Example

➢ Scenario: Multiple hospitals test different treatment plans (arms)

Treatment effectiveness (arm 1)		Observations	Rewards
Hospital 1:	+ ■ - □	■ ■ ■ ■ □ □ □	4/7
Hospital 2:	+ ■ - □	■ ■ □ □ □ □ □	2/7

True (global) reward: 3/7

➢ Problem: Biased data → Biased reward → Biased decision

➢ Main focus: Data heterogeneity

Federated Bandit

In Multi-Armed Bandit (MAB) settings, we consider:

- ➢ Heterogeneity
  - Bias in local learning
  - Problem may not be solved
- ➢ Decentralization
  - No central-controller
- ➢ Privacy
  - Protect agents' privacy in the worst cases during federation

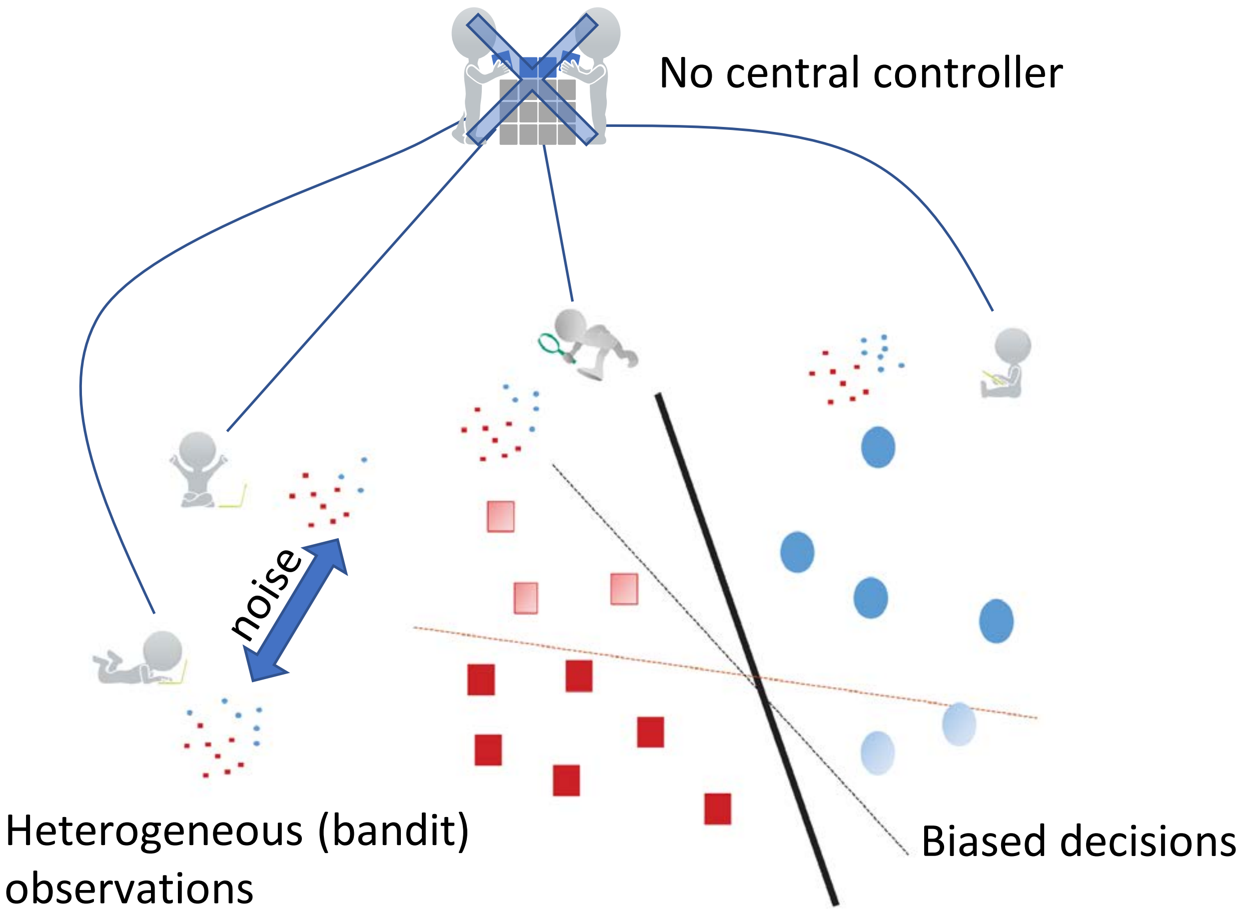

Road Map

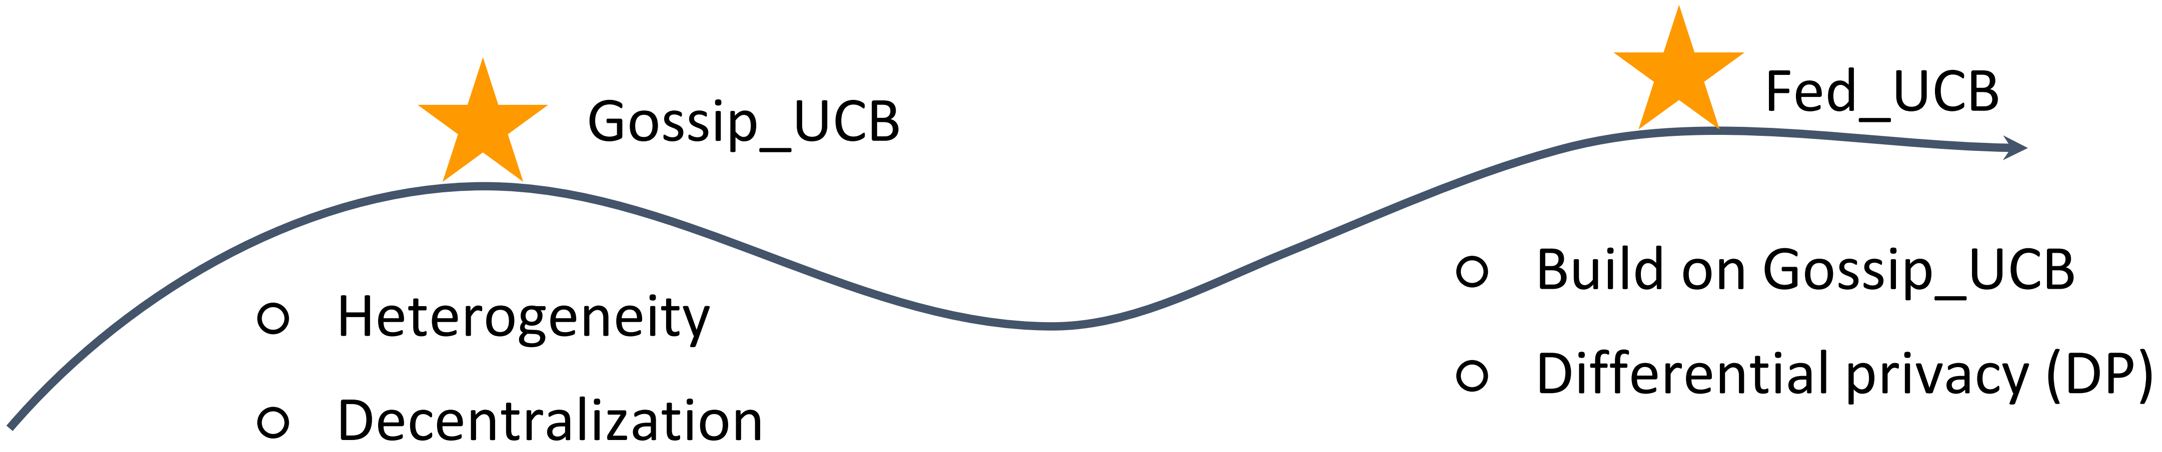
Gossip_UCB
Heterogeneity
Decentralization
Fed_UCB
Build on Gossip_UCB
Differential privacy (DP)

Regret

λ_2: 2nd-largest eigenvalue of the gossip matrix

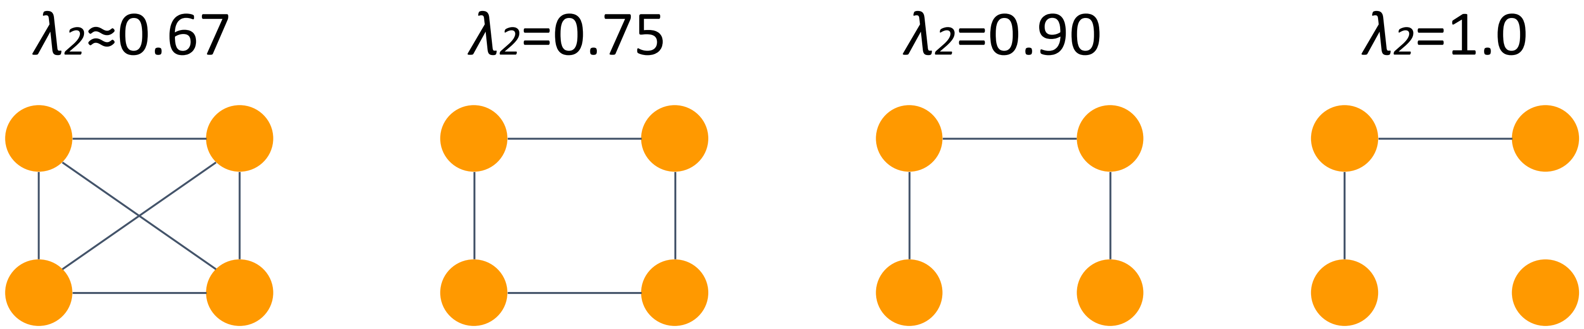

➢ Gossip_UCB (∝ log T, 1/connectivity)

$$O(\max\{\mathtt{poly}(N,M)\log T, \mathtt{poly}(N,M)\log_{\lambda_2^{-1}} N\})$$

➢ Fed_UCB (ε-DP, ∝ log T, 1/connectivity, 1/ε)

$$O(\max\{\frac{\mathtt{poly}(N,M)}{\epsilon}\log^{2.5} T, \mathtt{poly}(N,M)(\log_{\lambda_2^{-1}} N + \log T)\})$$

N: # agents *M*: # arms *T*: total time

Algorithm

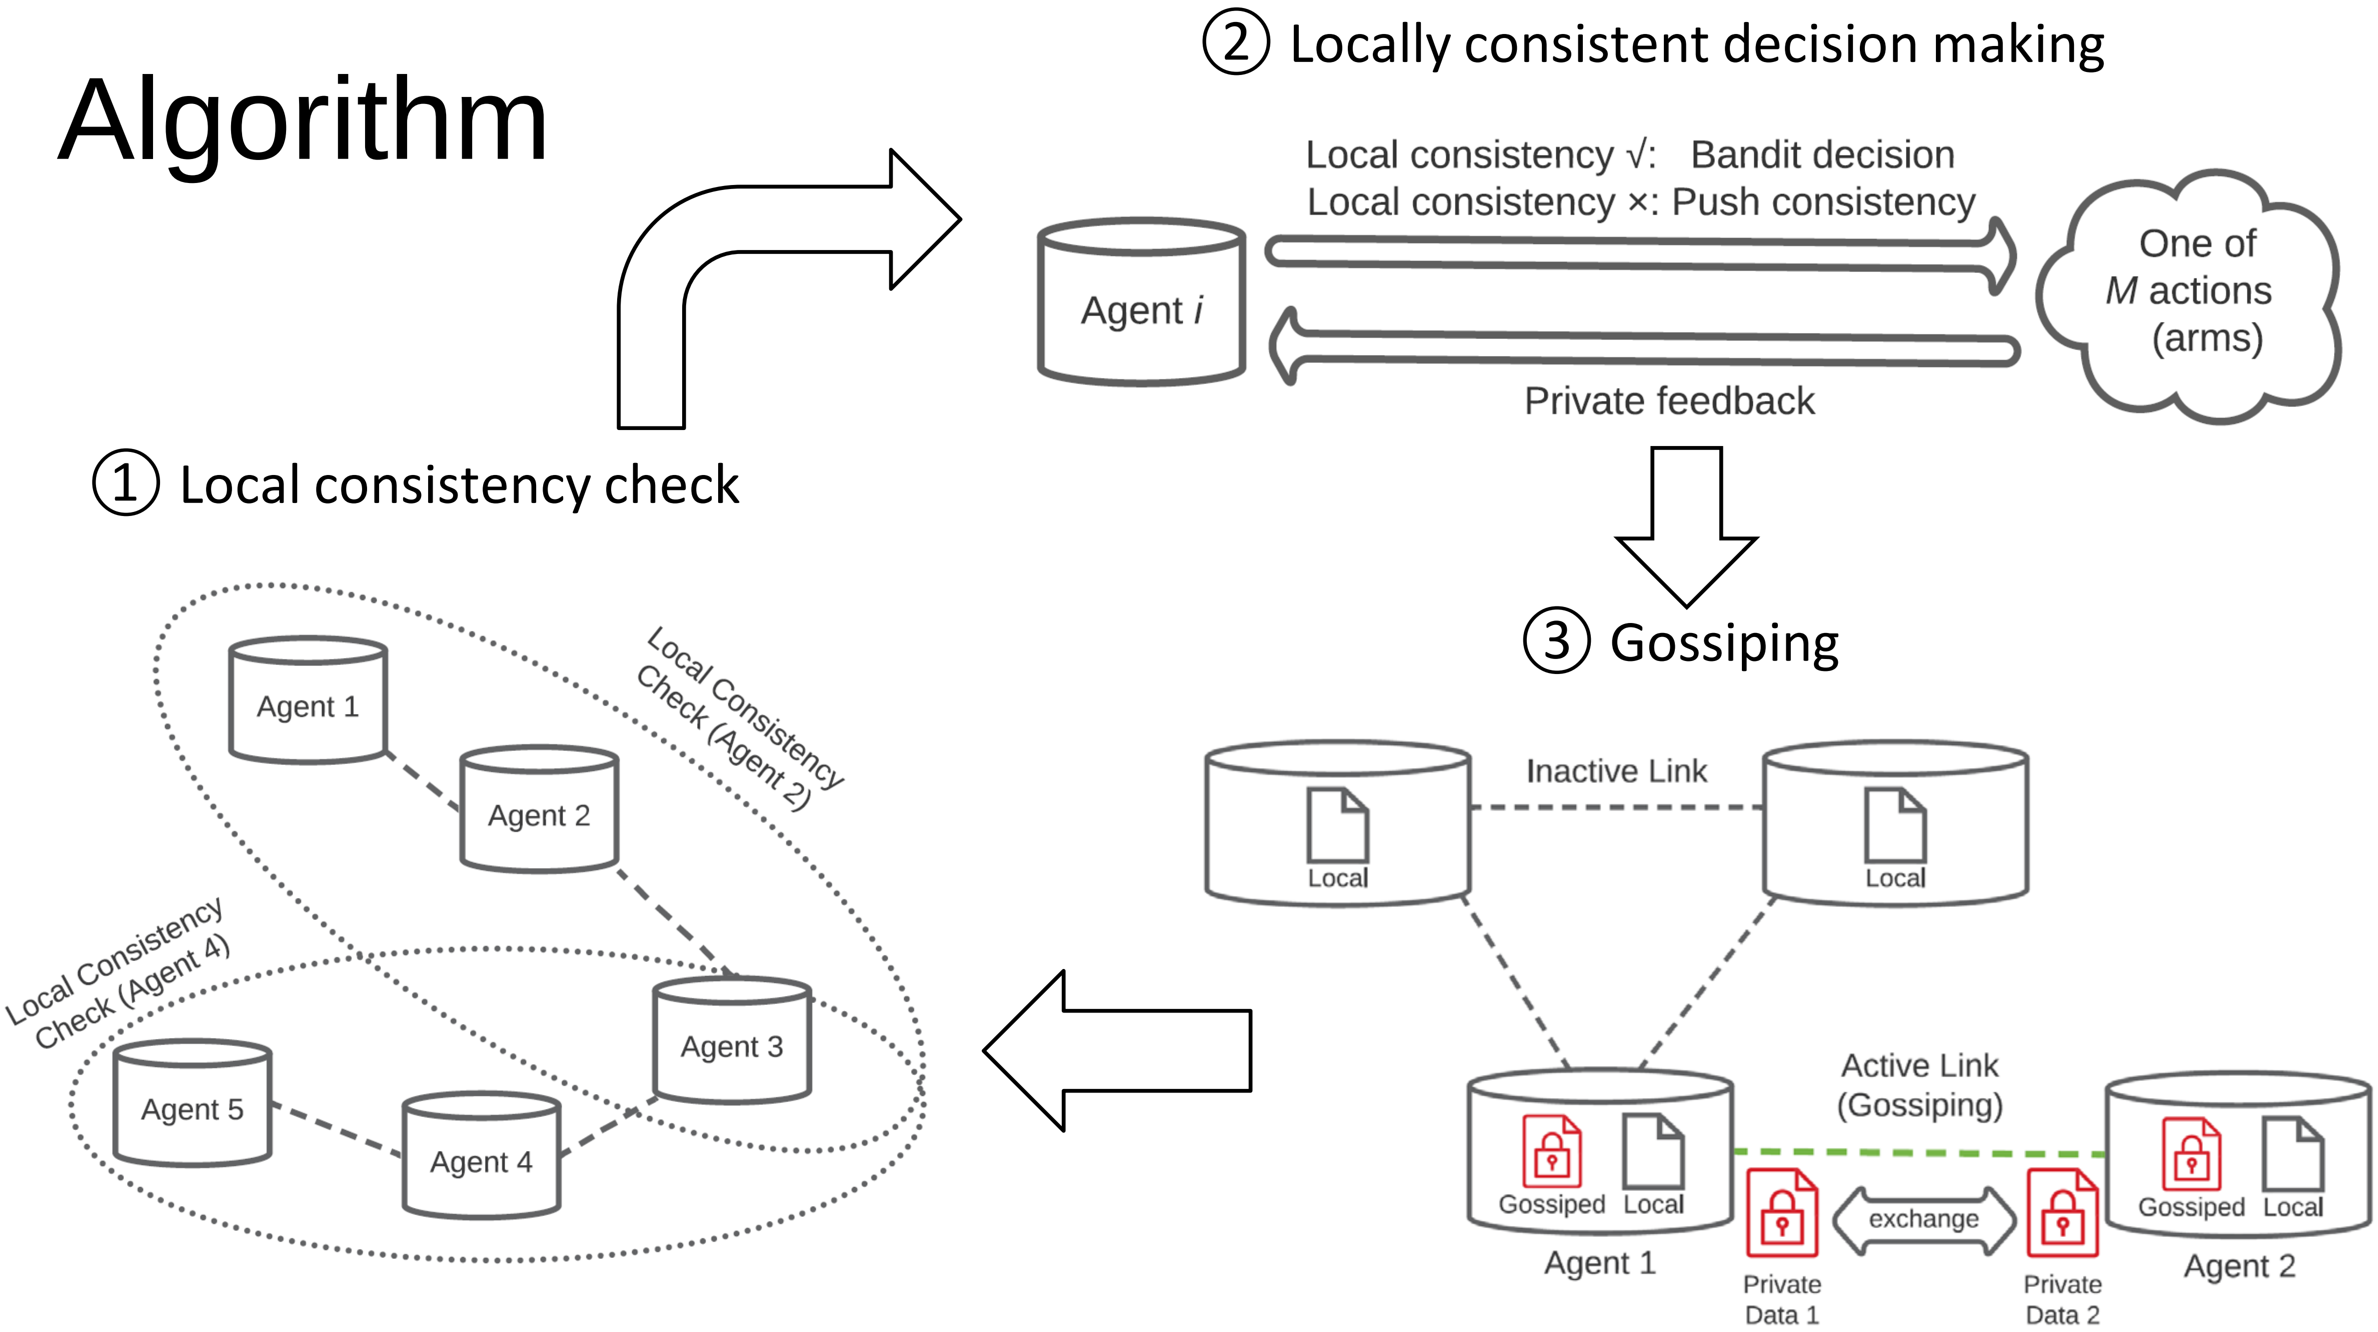
② Locally consistent decision making
Local consistency √: Bandit decision
Local consistency ×: Push consistency
Agent i
One of M actions (arms)
Private feedback
① Local consistency check
Local Consistency Check (Agent 2)
Local Consistency Check (Agent 4)
Agent 1
Agent 2
Agent 3
Agent 4
Agent 5
③ Gossiping
Inactive Link
Local
Local
Active Link (Gossiping)
Gossiped
Local
Agent 1
Private Data 1
exchange
Private Data 2
Gossiped
Local
Agent 2

Gossip_UCB

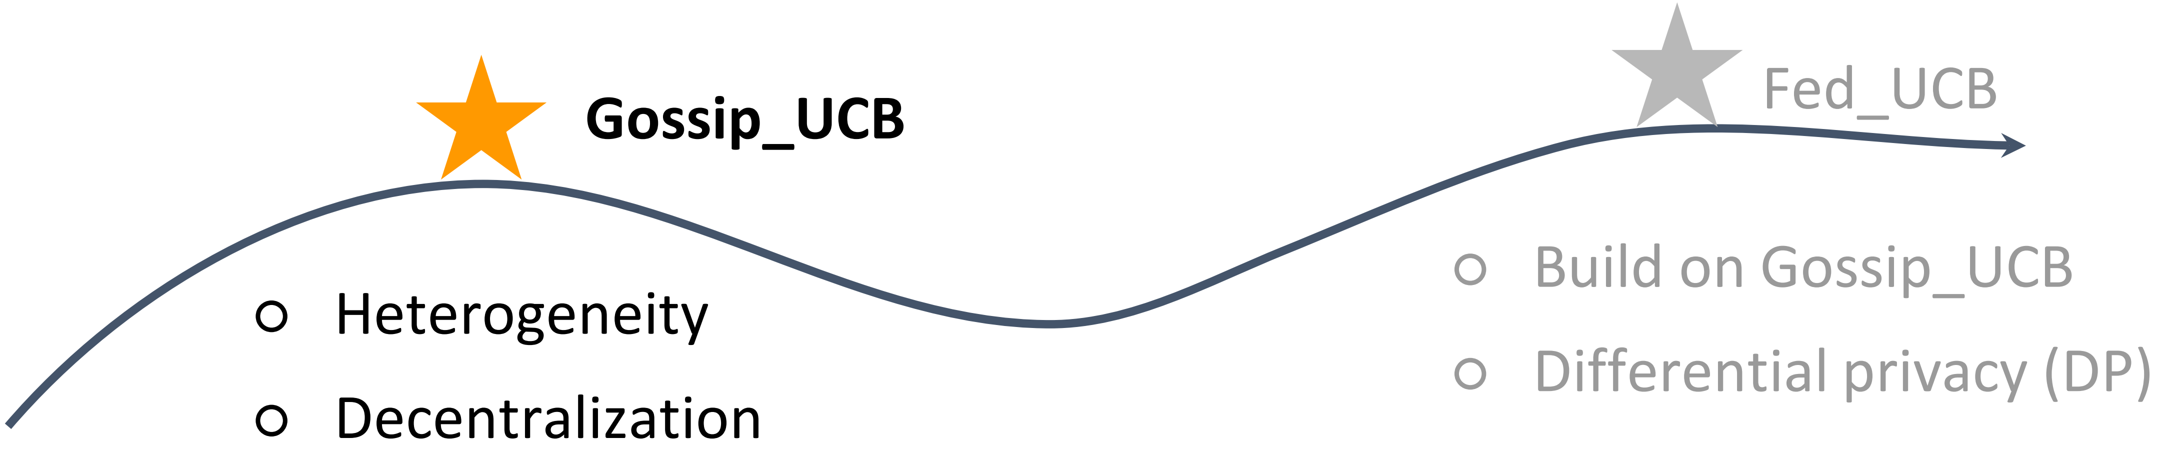
Gossip_UCB
Fed_UCB
Heterogeneity
Decentralization
Build on Gossip_UCB
Differential privacy (DP)

Gossip_UCB: Bandit Problem

- ➢ MAB with N agents, M arms (total time T):
  - ○ Agent i pulls arm $a_i(t)$ at time t, gets feedback $X_{i,a_i(t)}(t)$
  - ○ Ideal (global) feedback $X_{a_i(t)}(t)$
  - ○ Expectation of feedbacks: Local mean $\mu_{i,k}$
  - ○ Expectation of ideal feedbacks: Global mean μ_k
  - ○ Regret (μ_1 is the max):

$$R_i(T) = T\mu_1 - \sum_{t=1}^{T} \mathbb{E}\left[X_{a_i(t)}(t)\right]$$

- ➢ Homogenous setting:
  - ○ Agent gets exact feedback $\mu_{i,k} = \mu_k$
- ➢ Heterogenous setting:
  - ○ Agent gets biased/noisy feedback $\mu_{i,k} \neq \mu_k$

Gossip_UCB: Heterogeneous Feedbacks

- Global mean μ_k vs. local mean $\mu_{i,k}$:

$$\mu_k := \frac{1}{N}\sum_{i=1}^{N}\mu_{i,k}$$

- Estimates:
  - Sample mean (MEAN) $\tilde{X}_{i,k}(t)$: estimate of local mean (averaged observations)
  - Estimate of rewards (EST) $\vartheta_{i,k}(t)$: estimate of global mean
- Action: $a_i(t) = \arg\max_k \ \vartheta_{i,k}(t-1) + \boxed{C_{i,k}(t)}$ → Upper confidence bound (UCB)
- Heterogeneity:
  - Local mean ≠ Global mean → Sample mean ≠ Estimate of rewards

Design it!

Gossip_UCB: Gossiping

Communication among agents (classical gossiping [1]):

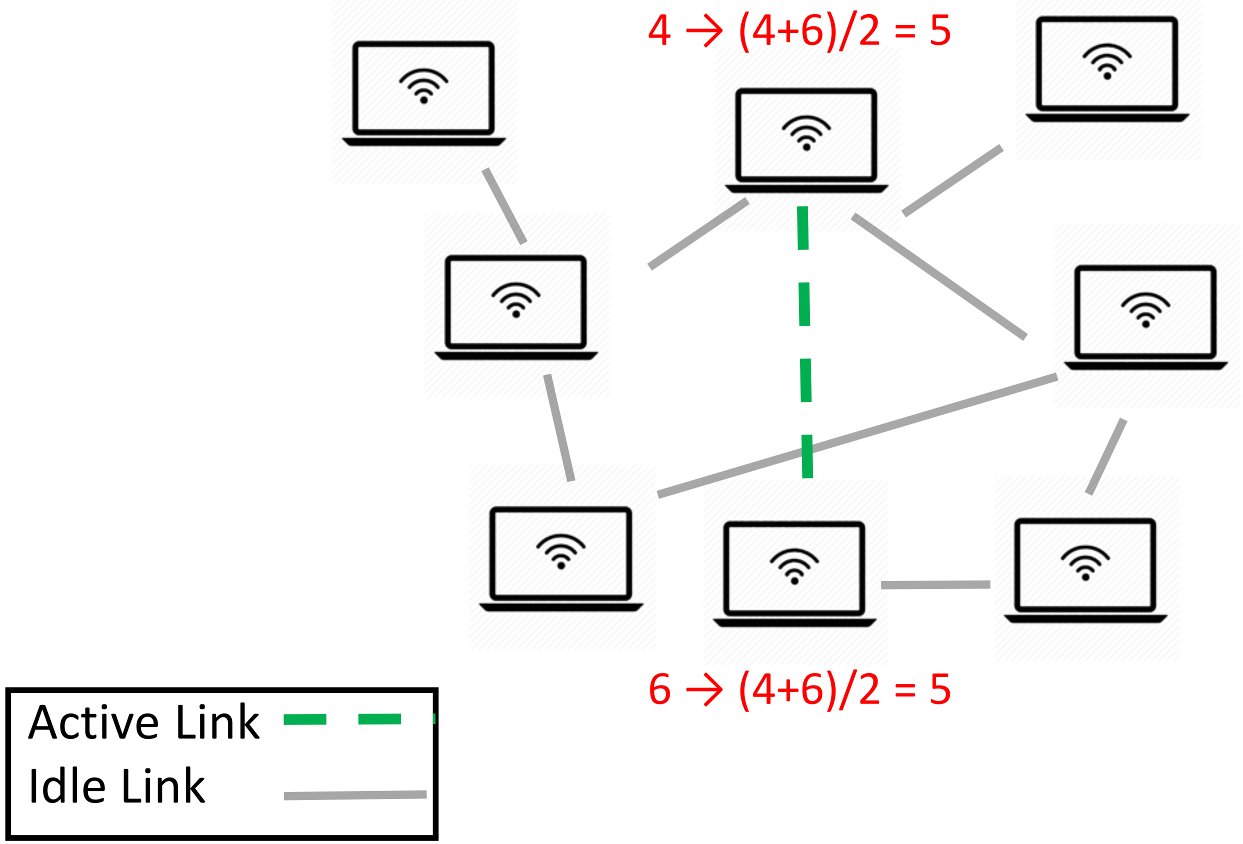

Features:

- ➢ One edge activated at each time *t*
- ➢ Selected agents on the edge exchange information
- ➢ Others do not update

GOSSIP = (ESTself + ESTother)/2

[1] Stephen Boyd, Arpita Ghosh, Balaji Prabhakar, and Devavrat Shah. 2006. Randomized Gossip Algorithms. *IEEE TIT* 52, 6 (2006), 2508–2530.

Gossip_UCB: Gossiping

Communication among agents (bandit gossiping [2]):

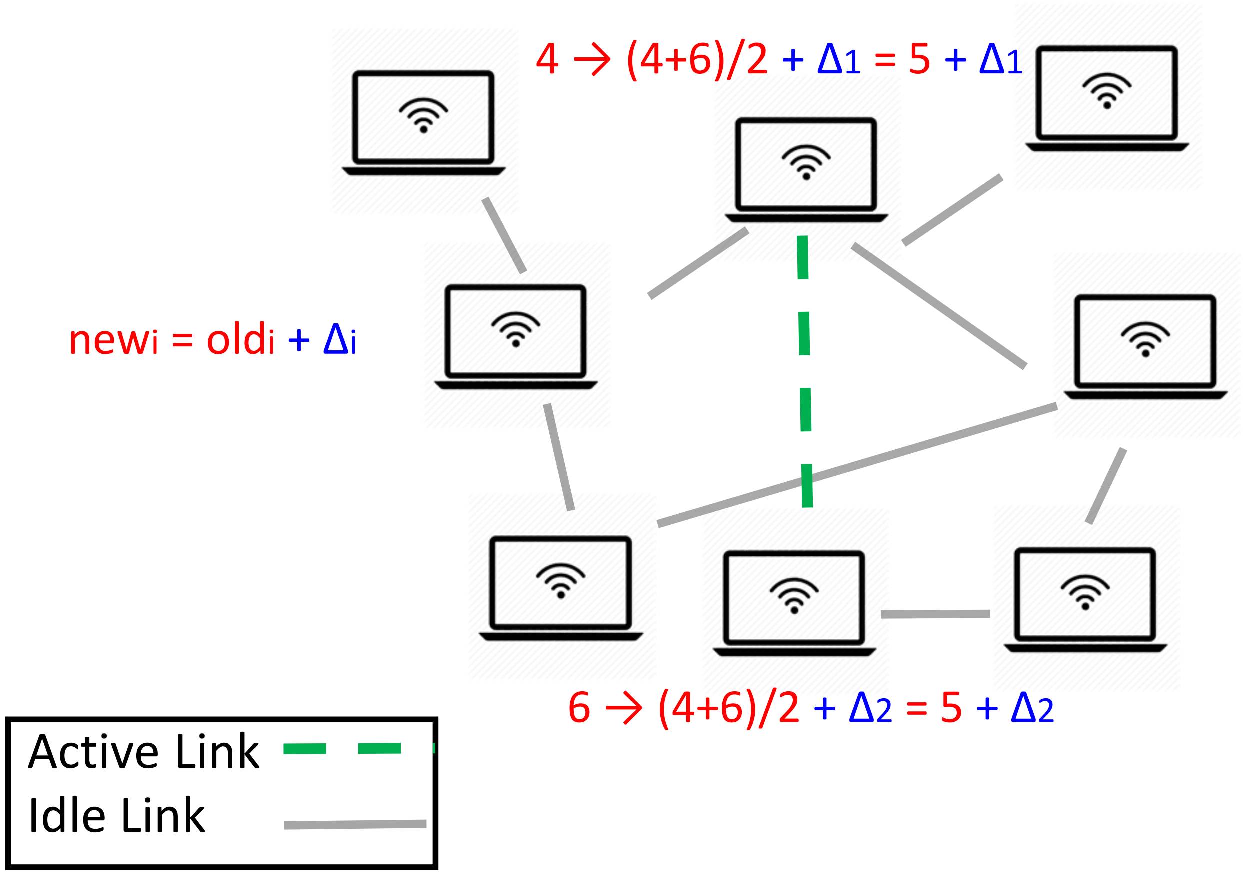

Features:

- ➢ One edge activated at each t
- ➢ Selected agents on the edge exchange information + gradient
- ➢ Others: gradient update

[2] Yang Liu, Ji Liu, and Tamer Başar. 2018. Differentially private gossip gradient descent. In *IEEE CDC*. IEEE, 2777–2782.

Gossip_UCB: Gossiping

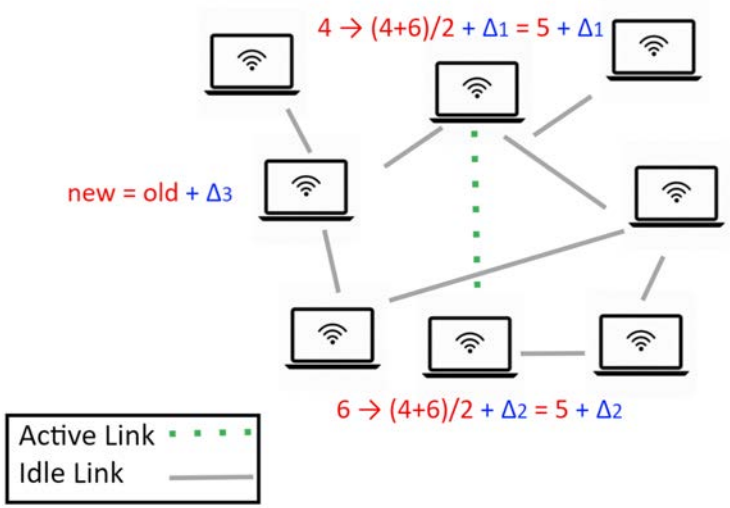

Communication among agents (bandit gossiping):

Each agent:

$$EST_t = \begin{cases} \boxed{GOSSIP} + MEAN_t - MEAN_{t-1} & \text{(Gossiping update)} \\ EST_{t-1} + \boxed{MEAN_t - MEAN_{t-1}} & \text{(Normal update)} \end{cases}$$

if *agent i is selected to gossip with agent j* **then**

  agent i sends $\vartheta_{i,k}(t-1)$ to agent j

  agent i receives $\vartheta_{j,k}(t-1)$ from agent j

  $\vartheta_{i,k}(t) := \boxed{\frac{\vartheta_{i,k}(t-1)+\vartheta_{j,k}(t-1)}{2}} + \tilde{X}_{i,k}(t) - \tilde{X}_{i,k}(t-1)$ *// gossiping update*

else

  $\vartheta_{i,k}(t) := \vartheta_{i,k}(t-1) + \boxed{\tilde{X}_{i,k}(t) - \tilde{X}_{i,k}(t-1)}$ *// normal update*

end

Local information (#pulls) sharing:

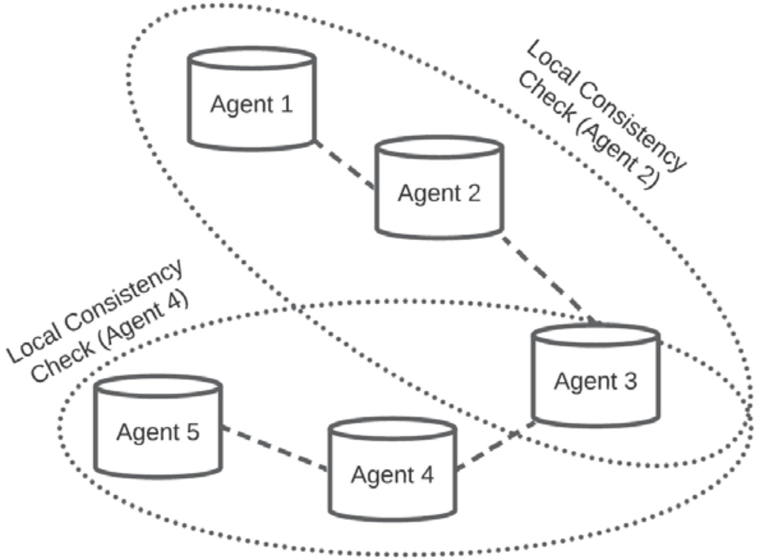

Gossip_UCB: Consistency

- ➢ Problem:
  - ○ Arm true estimates depend on global estimates (average of local estimates)
  - ○ #Pull affects the calculation of UCB

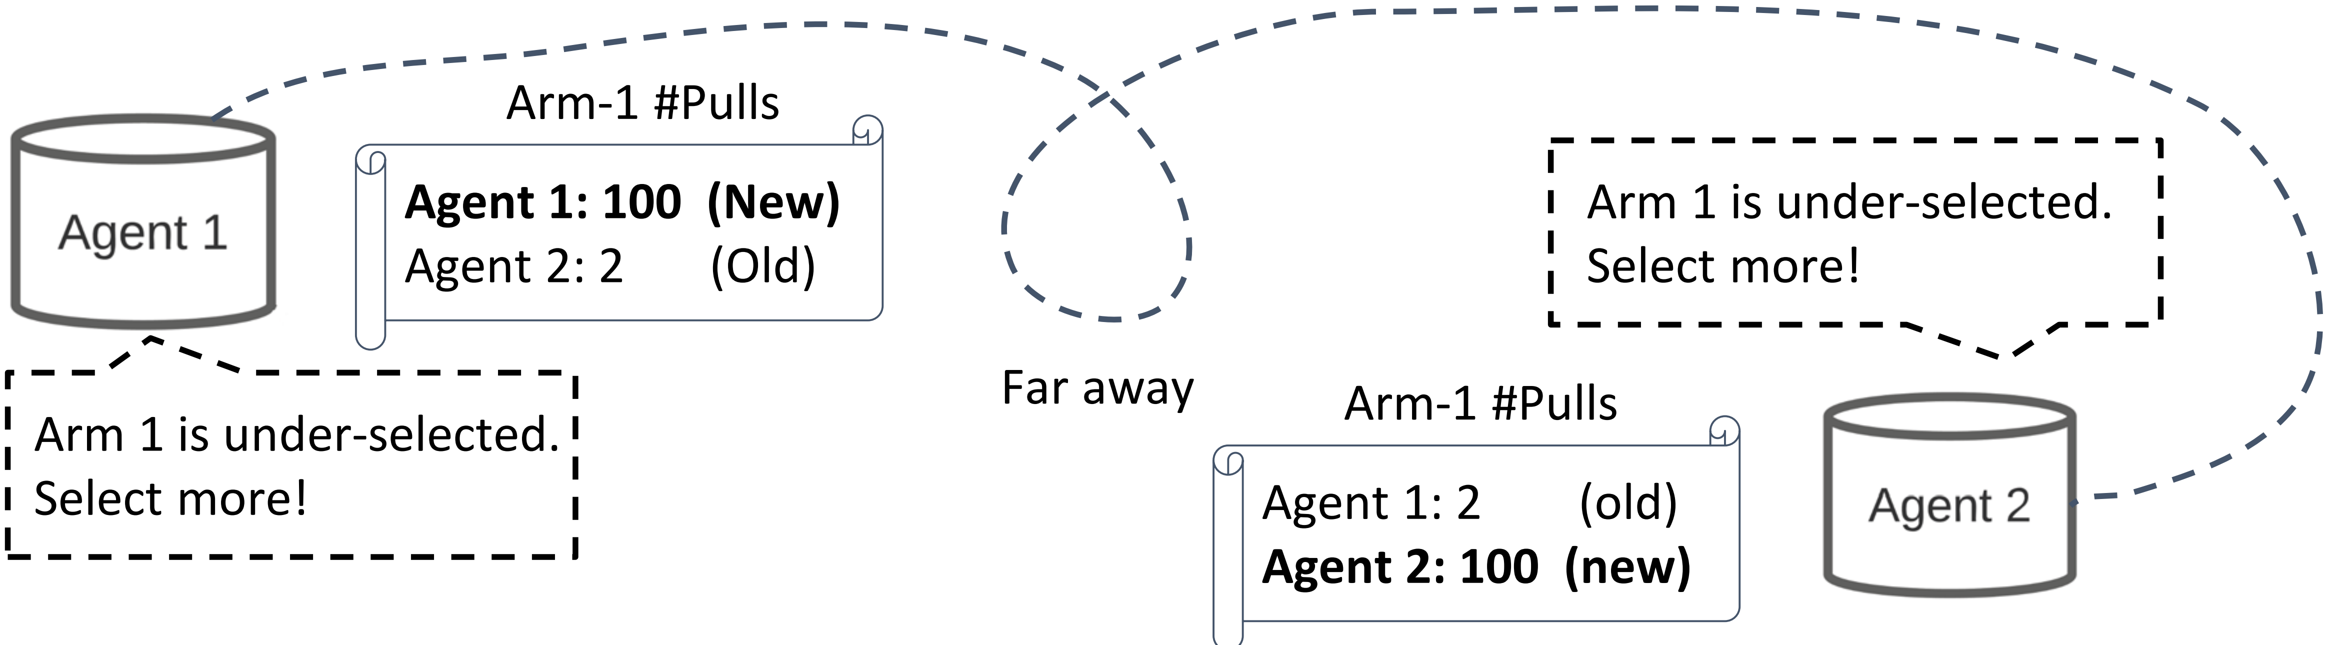

Gossip_UCB: Consistency

- Local consistency:
  - #LocalPulls is close to the local estimate of #GlobalPulls (Lemma 2)
  - Local estimate is not a bottleneck

- Global consistency:
  - Max #LocalPulls ≤ 2 × #LocalPulls$_i$, for all agent i (Lemma 3)

- Summary:
  - Bound the information inconsistency due to propagation delay
  - Facilitate a fully-decentralized solution

Gossip_UCB: Concentration Bound

- Concentration Bound for Local Estimates (Theorem 1)
  - With some conditions and a high probability ($1 - p_0$):

$$\mathbb{P}(|\boxed{\vartheta_{i,k}(t)} - \boxed{\mu_k}| \geq C_{i,k}(t)) < \frac{2}{t^2}$$

where

$$C_{i,k}(t) = \sqrt{\frac{2N}{\boxed{n_{i,k}(t)}} \log t + \alpha_1}$$

  - Notations: local estimates, global mean, #LocalPulls (global consistency)
  - Challenges:
    - Coupling effects of gossiping and bandit learning

Gossip_UCB: Proof Overview

- ➢ Guarantees on the Consistency among Agents
  - Information propagation (Lemma 1)
  - Actual local consistency (Lemma 2)
  - Global consistency (Lemma 3)
- ➢ Concentration Bound for Local Estimates
  - Bound |*LocalEst* - E[*LocalEst*]| (Lemma 4)
  - Bound |E[*LocalEst*] - *GlobalMean*| (Lemma 5)
- ➢ Regret Upper Bound for Gossip_UCB (Theorem 1)

Fed_UCB

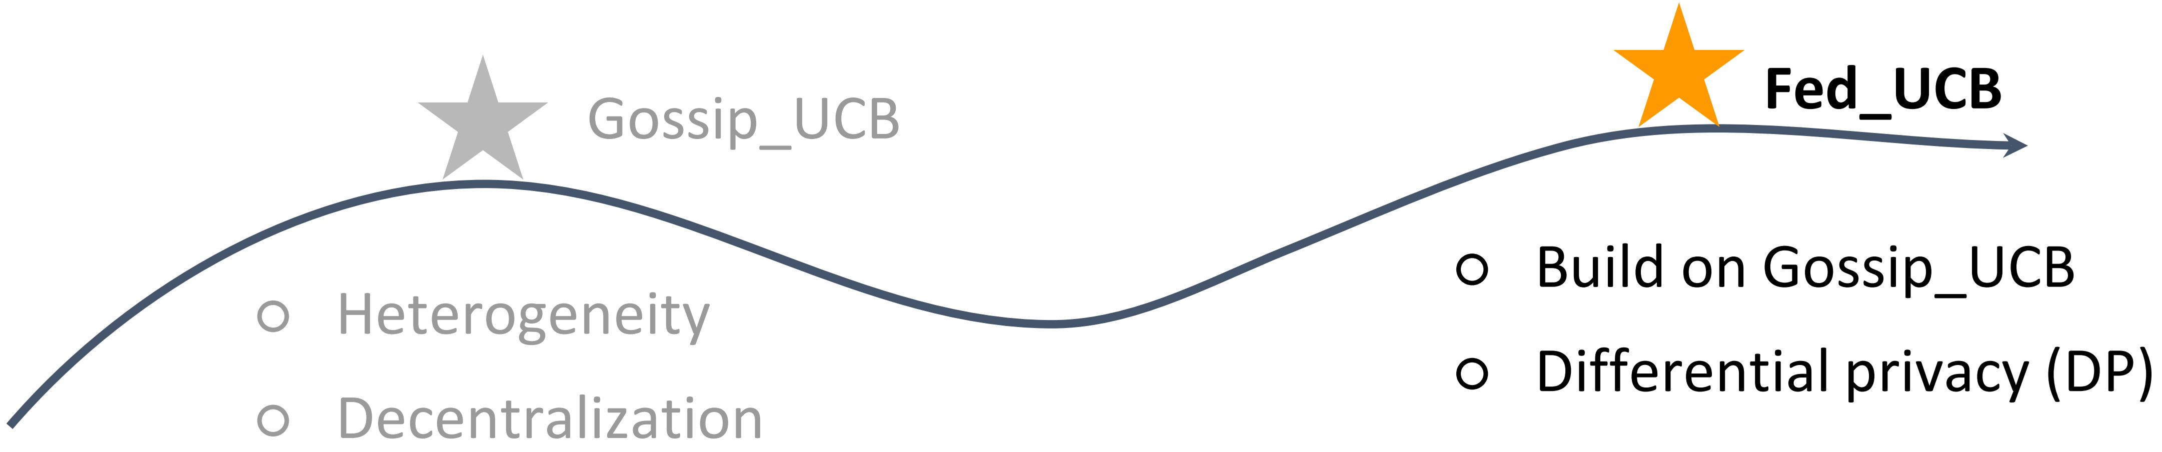
Gossip_UCB
Heterogeneity
Decentralization
Fed_UCB
Build on Gossip_UCB
Differential privacy (DP)

Fed_UCB: Differential Privacy

- ➢ Why necessary?
  - ○ Directly leaking some information that might appear to be "anonymized" can be used to cross-reference with other datasets to breach privacy [3]
  - ○ Worst-case privacy guarantee
- ➢ Differential privacy (DP) [4]:

A (randomized) algorithm $\mathcal{B}$ is ϵ-differentially private if for any adjacent streams $\{X_{i,k}(t)\}_{t=1}^{T}$ and $\{X'_{i,k}(t)\}_{t=1}^{T}$, and for all sets $\mathcal{O} \in \mathcal{C}$,

$$\mathbb{P}\left[\mathcal{B}(\{X_{i,k}(t)\}_{t=1}^{T}) \in \mathcal{O}\right] \leq e^{\epsilon} \cdot \mathbb{P}\left[\mathcal{B}(\{X'_{i,k}(t)\}_{t=1}^{T}) \in \mathcal{O}\right].$$

[3] Latanya Sweeney. 2000. Simple demographics often identify people uniquely. *Health (San Francisco)* 671, 2000, 1–34.
[4] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating noise to sensitivity in private data analysis. *In Theory of cryptography conference*. Springer, 265–284.

Fed_UCB: Online DP

- ➢ Naive method:
  - ○ Add Laplacian noise Lap(T/ε) to each observation
- ➢ Partial sum [5]:
  - ○ Add Laplacian noise Lap((log T)/ε) following a tree structure

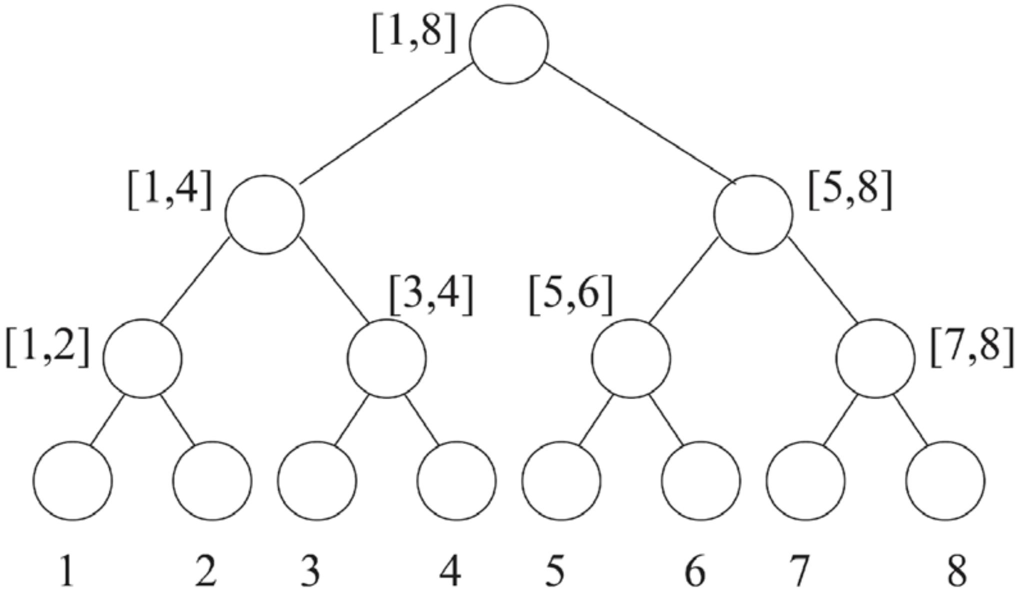

Example:

$$\text{Node 4: Noise}_{[1,4]}$$
$$\text{Node 6: Noise}_{[1,4]} + \text{Noise}_{[5,6]}$$
$$\text{Node 7: Noise}_{[1,4]} + \text{Noise}_{[5,6]} + \text{Noise}_{7}$$

[5] T-H Hubert Chan, Elaine Shi, and Dawn Song. 2011. Private and continual release of statistics. *ACM Transactions on Information and System Security* 14, 3 (2011), 26.

Fed_UCB: Concentration Bound

- Concentration Bound for Local Estimates
  - Guarantee ε-DP, with some conditions and a high probability ($1 - \boxed{\frac{2N}{n_{i,k}(t)}} - p_0$):

$$\mathbb{P}(|\tilde{\vartheta}_{i,k}(t) - \mu_k| \geq \tilde{C}_{i,k}(t)) < \frac{2}{t^2}$$

  where

$$\tilde{C}_{i,k}(t) = \alpha_1 + \sqrt{2N\left(\boxed{\frac{128N\log^2 T \cdot \log t \cdot \log n_{i,k}(t)}{n_{i,k}^2(t)\epsilon^2}} + \frac{1}{n_{i,k}(t)}\right)\log t}.$$

- Compared with Gossip_UCB: two changes
  - Upper confidence bound
  - Probability

Experiments

- ➢ Synthesize Fig. (a), Fig. (b)
- ➢ UCI Fig. (c)

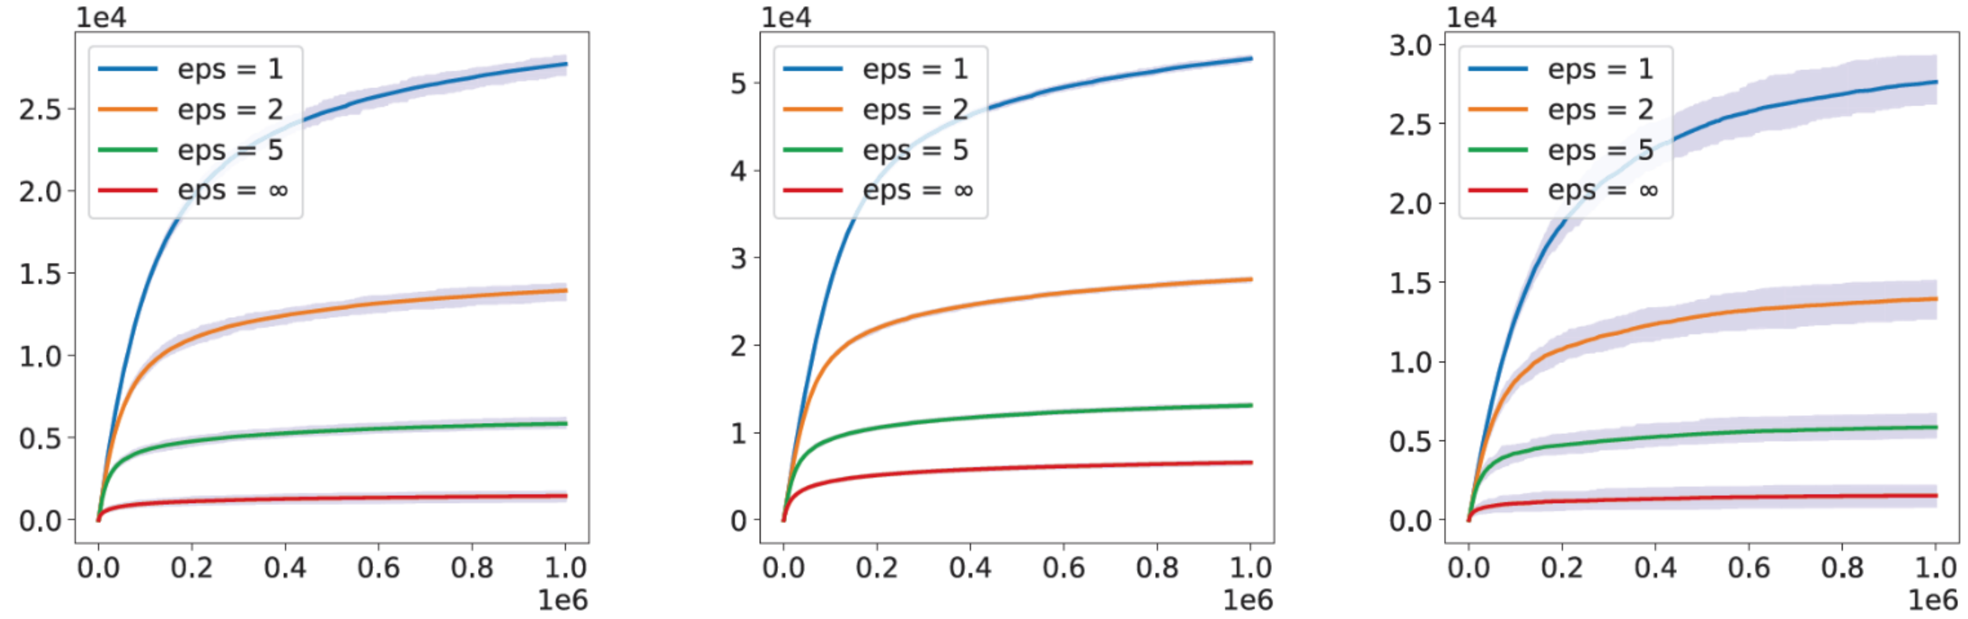

(a) (b) (c)

Thanks for your attention!

ACKNOWLEDGEMENT

This work is partially supported by the National Science Foundation (NSF) under grant IIS-2007951 and the Office of Naval Research under grant N00014-20-1-22