JPMorganChase

DMM: Distributed Matrix Mechanism for Differentially-Private Federated Learning Based on Constant-Overhead Linear Secret Resharing

Alexander Bienstock, Ujjwal Kumar, Antigoni Polychroniadou

JPMorgan AI Research & JPMorgan AlgoCrypt CoE

Federated Learning

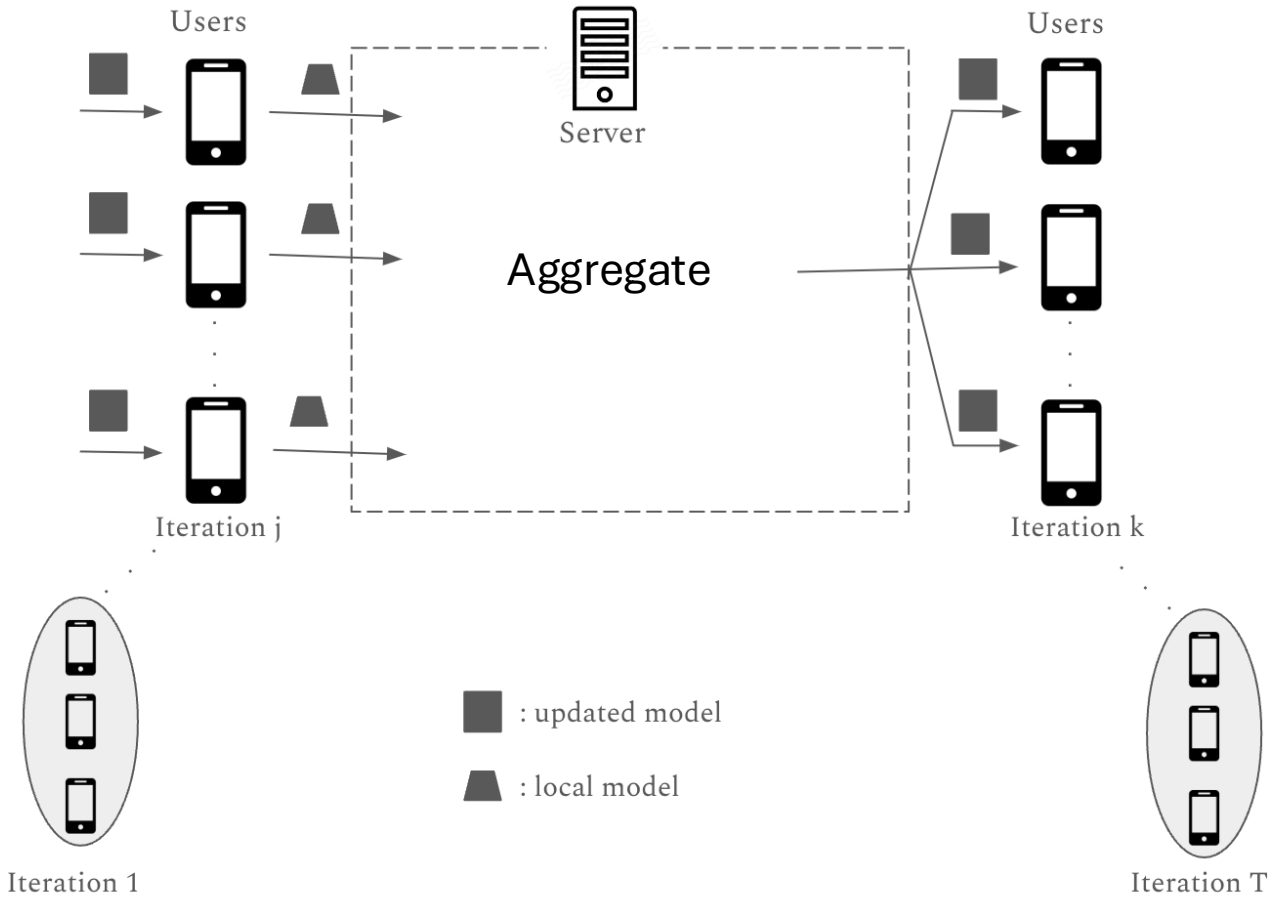

Lots of attention in academia and industry recently:

- Gboard and Apple keyboard next-word suggestion models
- Apple voice assistant training

Differential Privacy (DP)

- Needed for many applications:
  - People may type their SSN or other PII into their phones
  - Language Models *memorize* training data (also applicable for other models)

- (ε, δ)-DP for mechanism M: for any datasets D, D' differing by one user and any measurable outcome S:
  - $\Pr[M(D) \in S] \leq e^{\varepsilon} \cdot \Pr[M(D') \in S] + \delta$

Central Differential Privacy

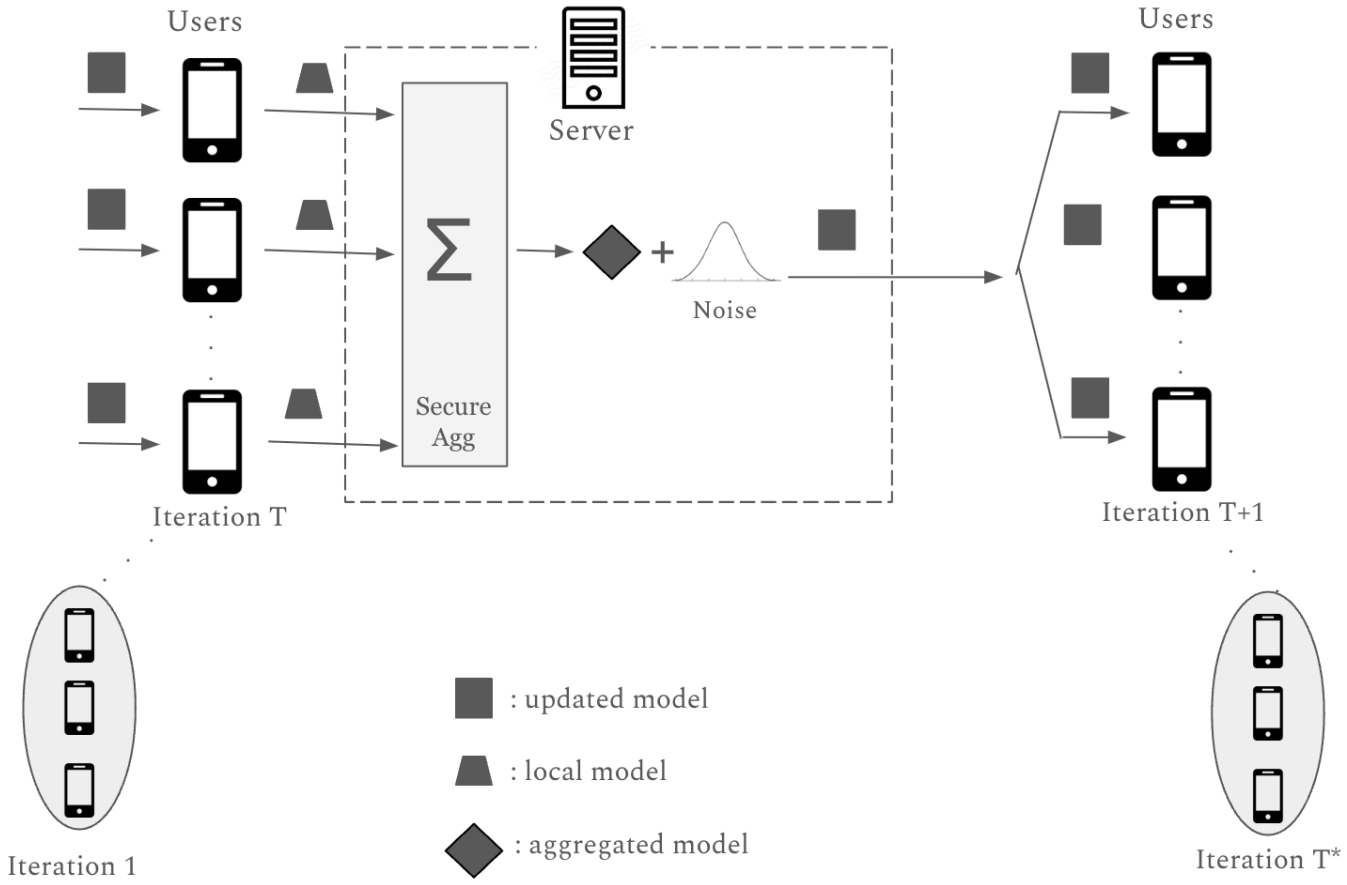

No DP with respect to server; trust server to add DP

Server can reuse/correlate noise across iterations !

- (Matrix Mechanism [CSS10])
- Yields better accuracy

Distributed Differential Privacy

Distributed Discrete Gaussian Mechanism [KLS21]

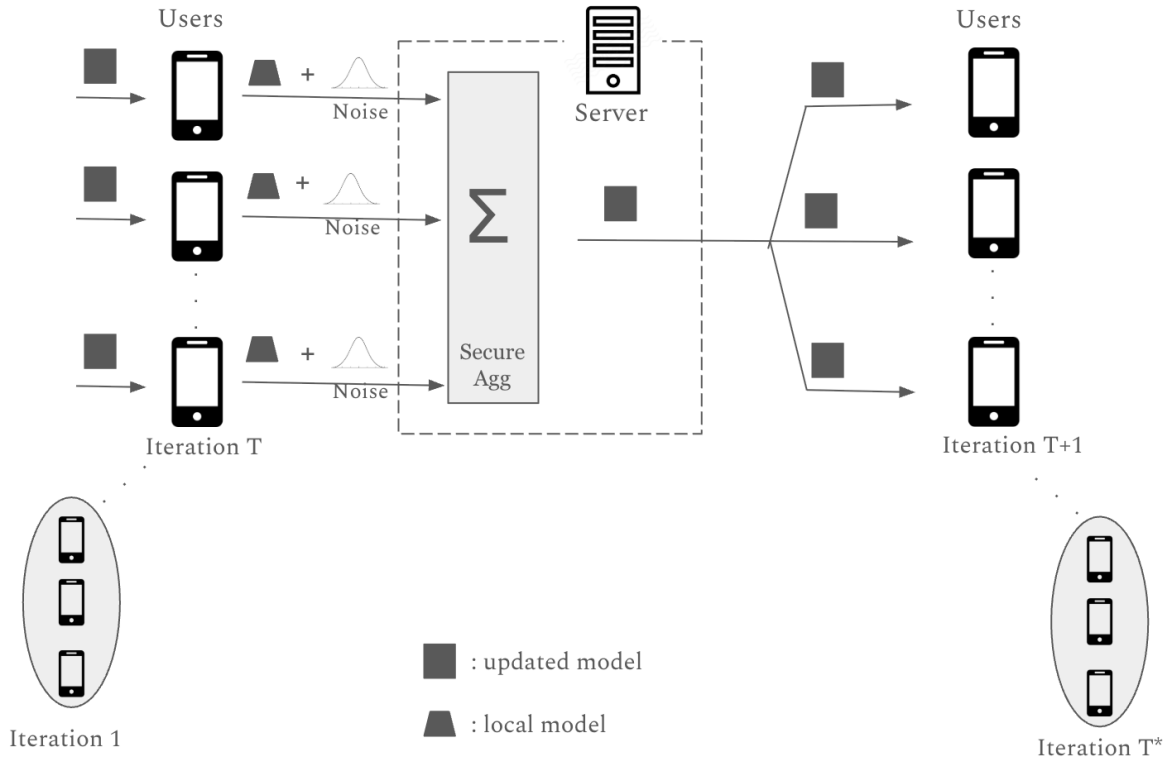

- DP with respect to everyone (including server)
- Noise is not correlated across rounds ☹
  - Leads to worse accuracy

Our Contribution: Distributed Matrix Mechanism (DMM)

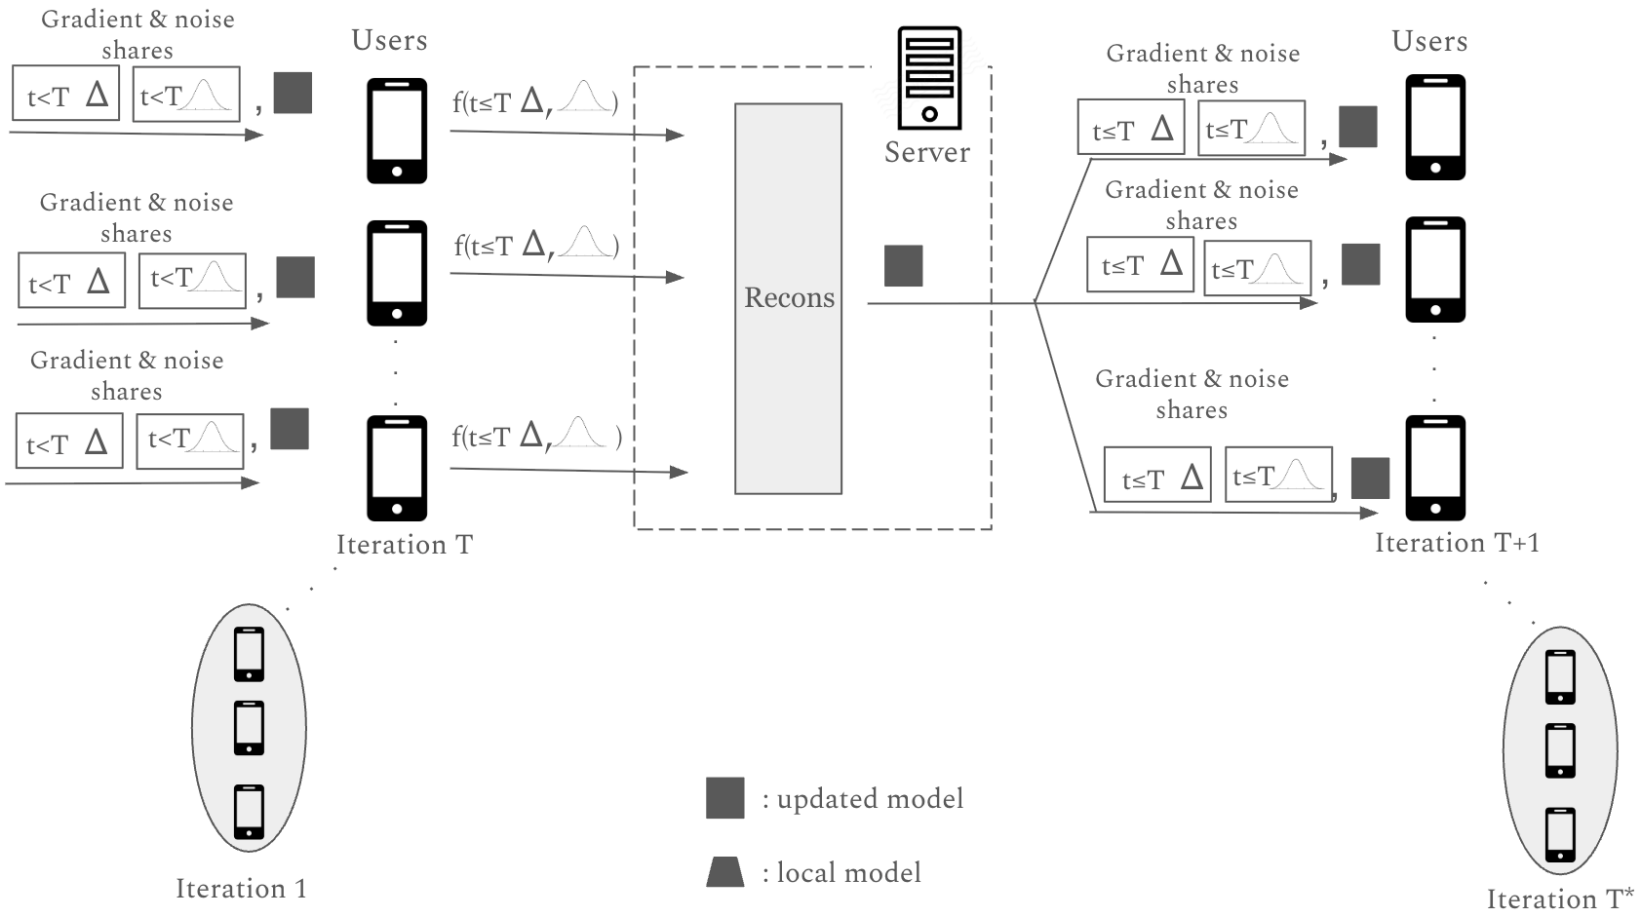

- Local DP
- And correlated noise!
- Using novel, efficient instantiation of cryptographic "secret (re)sharing"

Our Contribution: Distributed Matrix Mechanism (DMM)

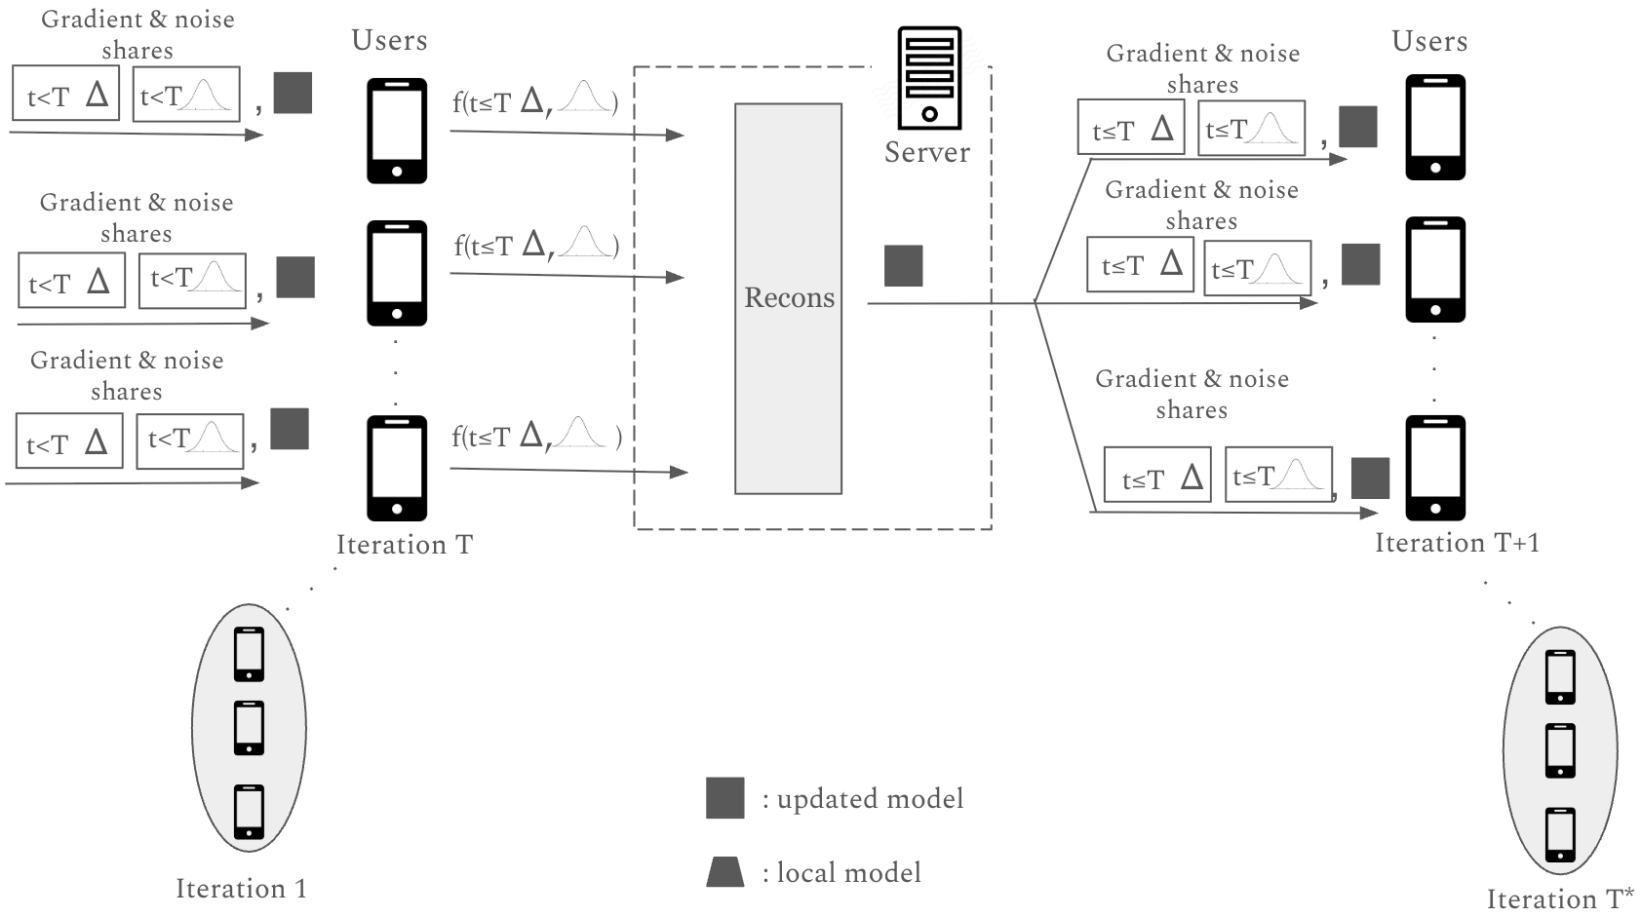

How to share?

- Baseline expensive: n^2 overhead per secret per round

Constant-Overhead Packed Resharing Protocol

Our technique: $O(1)$ overhead per secret per round!

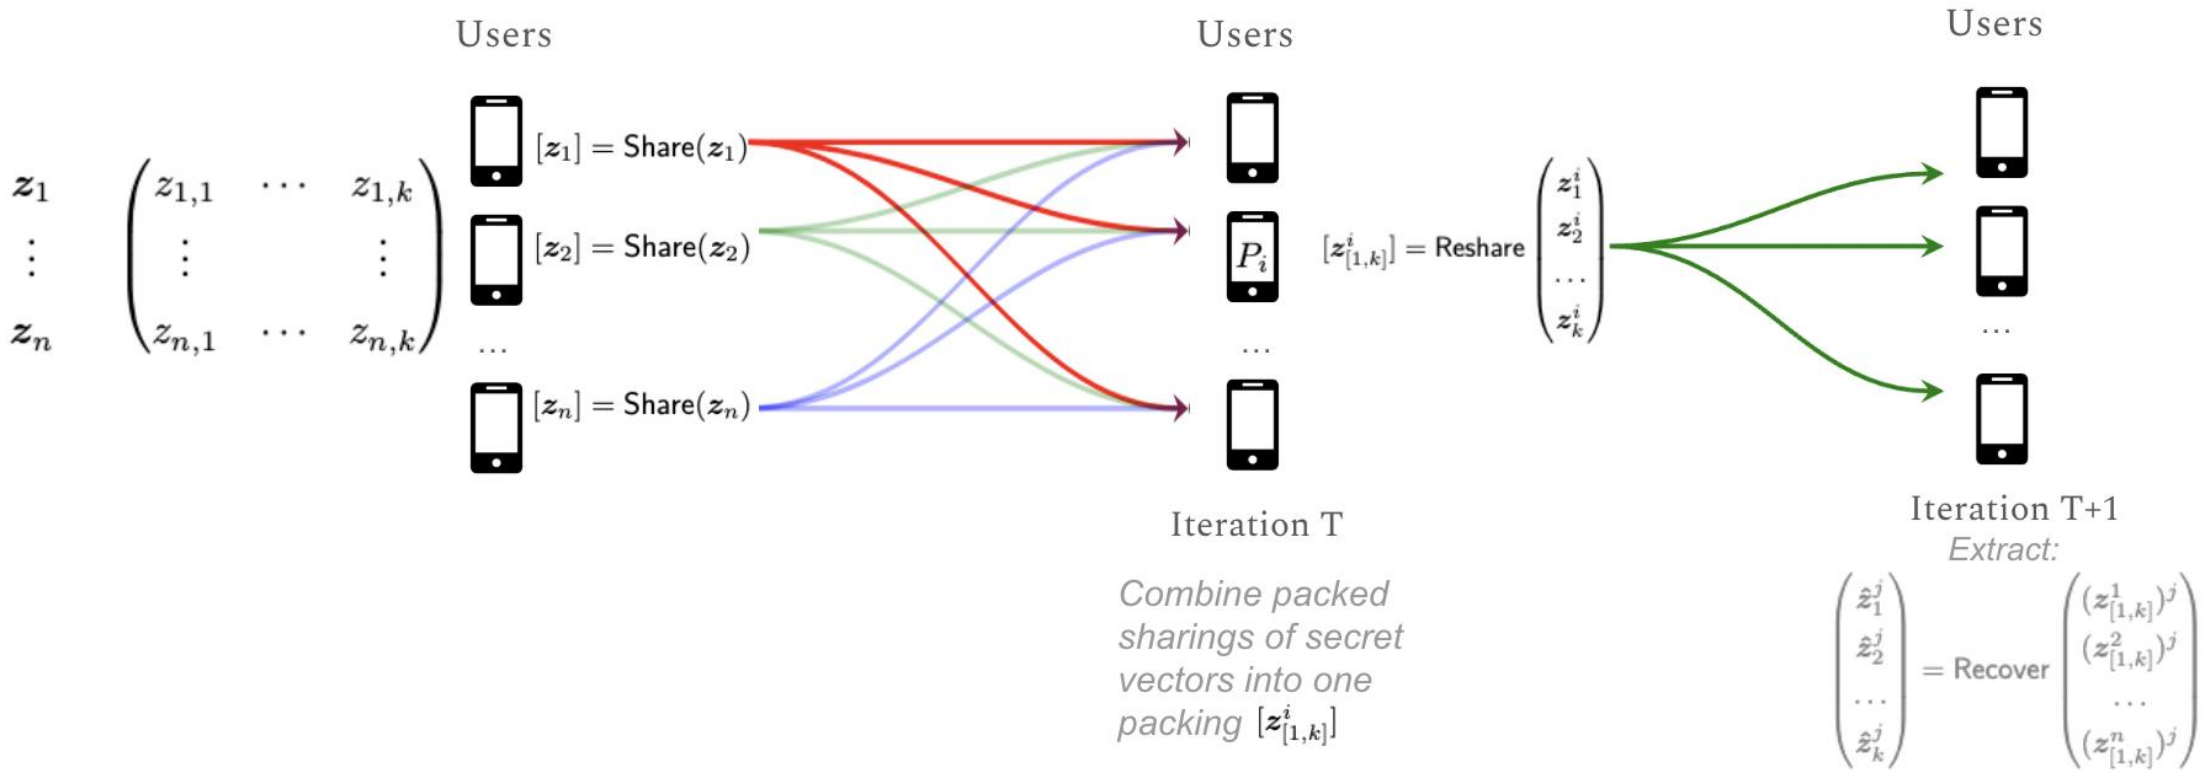

Experiments: Accuracy

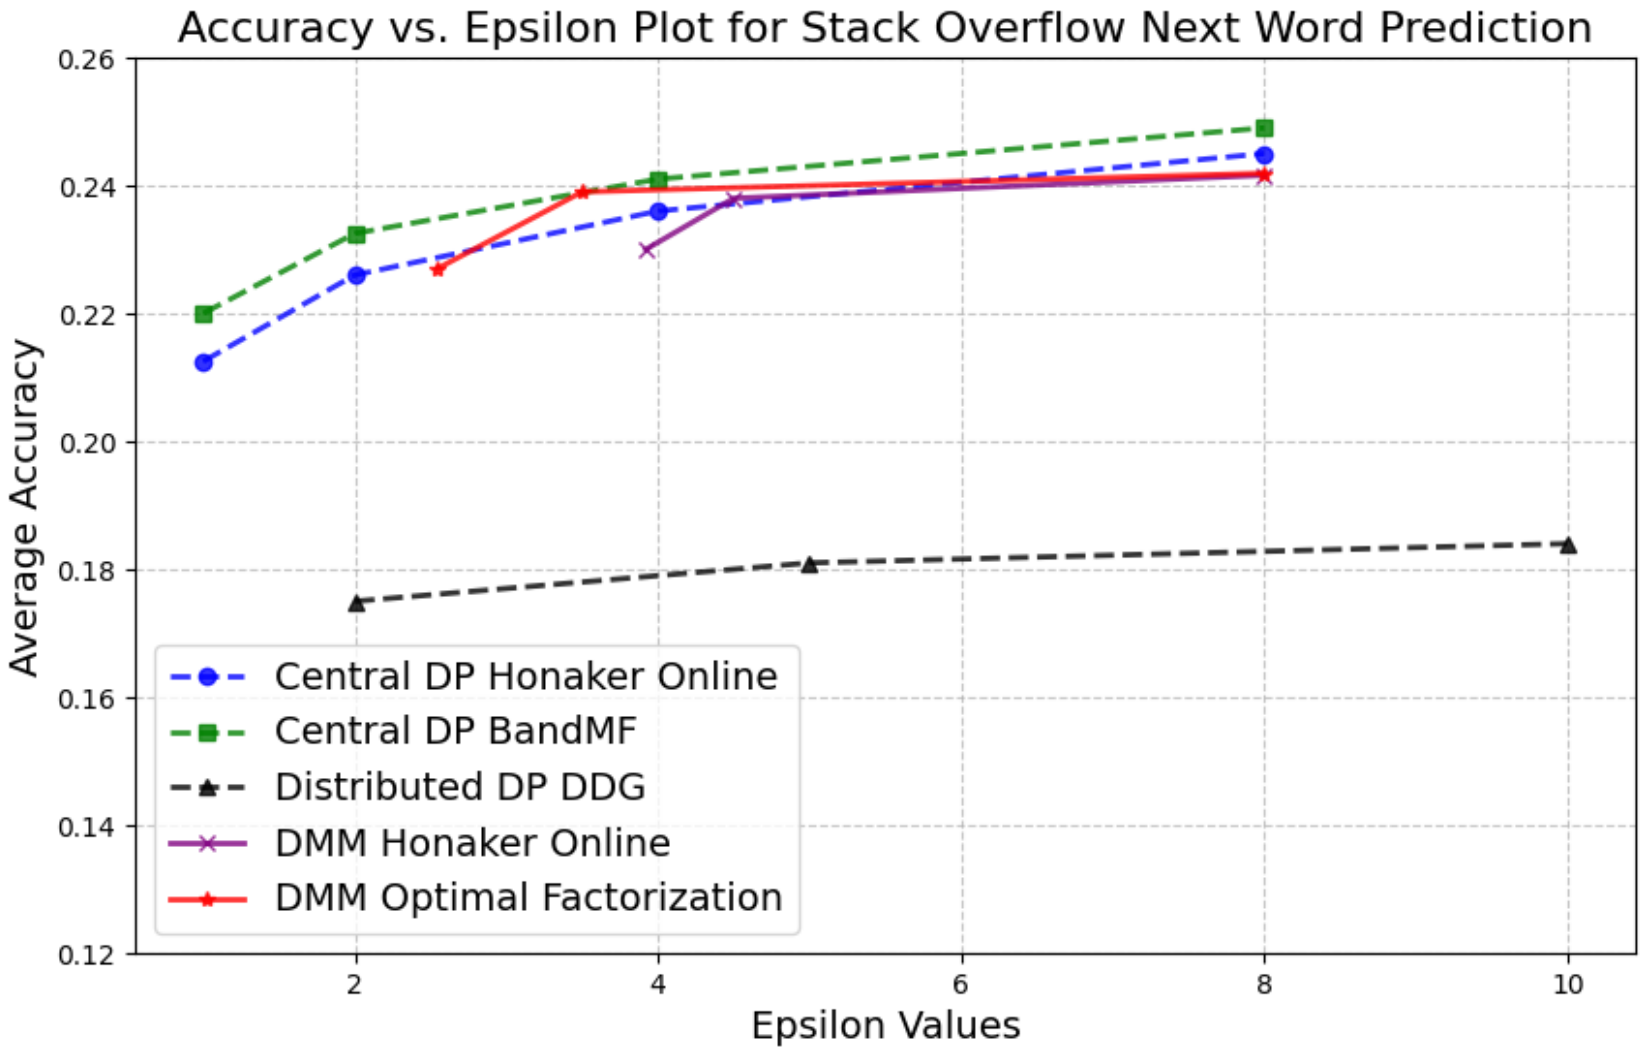

Thanks!