RÉPUBLIQUE FRANÇAISE
Liberté
Égalité
Fraternité

FRANCE 2030 PROGRAMME DE RECHERCHE INTELLIGENCE ARTIFICIELLE

cnrs

GRANITE: a Byzantine-Resilient Dynamic Gossip Learning Framework

Workshop on Adversarial Threats on Real Life Learning Systems

17/09/2025

Yacine Belal, Mohamed Maouche, **Sonia Ben Mokhtar**, Anthony Simonet-Boulogne

Federated Learning [MCM17]

1. **Model Broadcasts:** Server sends global model θ^t to all users $N = \{1,2,\dots,n\}$
2. **Local Training:** Each user i optimizes locally
$$\theta_i^t = \theta^t - \eta \nabla L(\theta^t; D_i)$$
3. **Model Upload:** Users return updated models θ_i^t to the server
4. **Model Aggregation:** Server aggregates client models
$$\theta^{t+1} = \frac{1}{\sum_{i \in N} |D_i|} \sum_{i \in N} |D_i| \theta_i^t$$

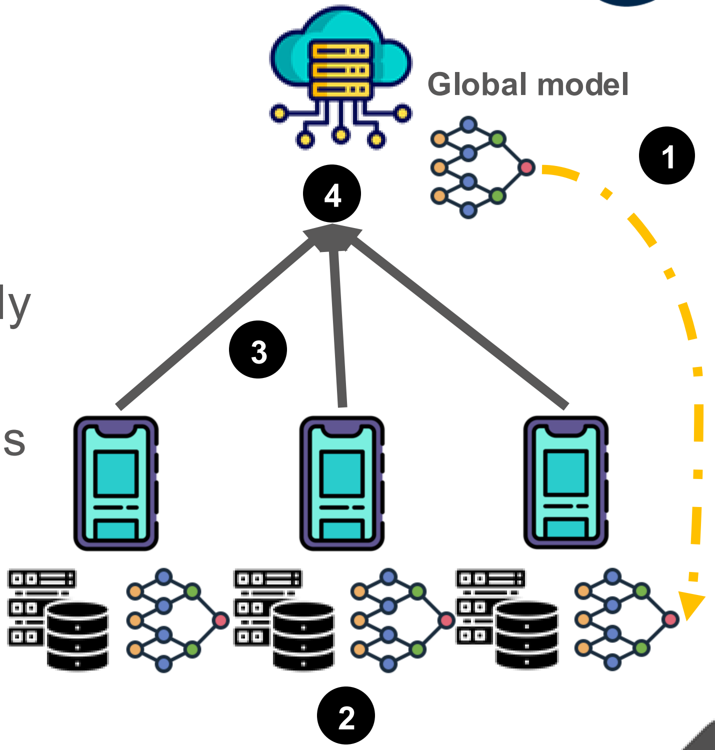

[MCM17] McMahan et al., **Communication-efficient learning of deep networks from decentralized data**. AISTATS'17.

Federated Learning [MCM17]

- **Single point of failure** [KAI21]

The central server's critical role makes the system vulnerable to failure and attacks

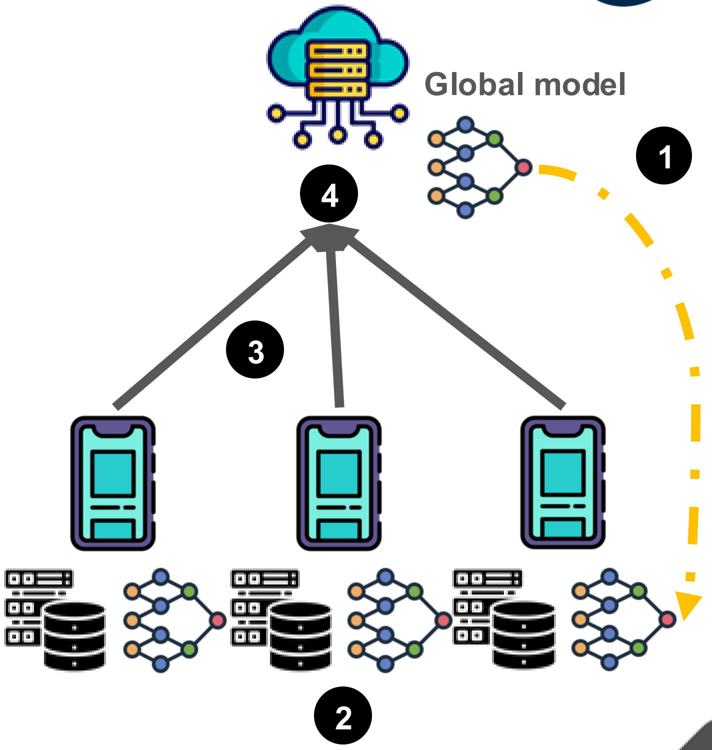

[MCM17] McMahan et al., **Communication-efficient learning of deep networks from decentralized data**. AISTATS'17.
[KAI21] Kairouz et al., **Advances and open problems in federated learning.** Fondations and Trends in Machine Learning'21.
.
.
.

Federated Learning [MCM17]

- **Single point of failure** [KAI21]

The central server's critical role makes the system vulnerable to failure and attacks

- **Governance drawbacks**

Power monopoly [VAN24]
Lack of transparency [GU24]

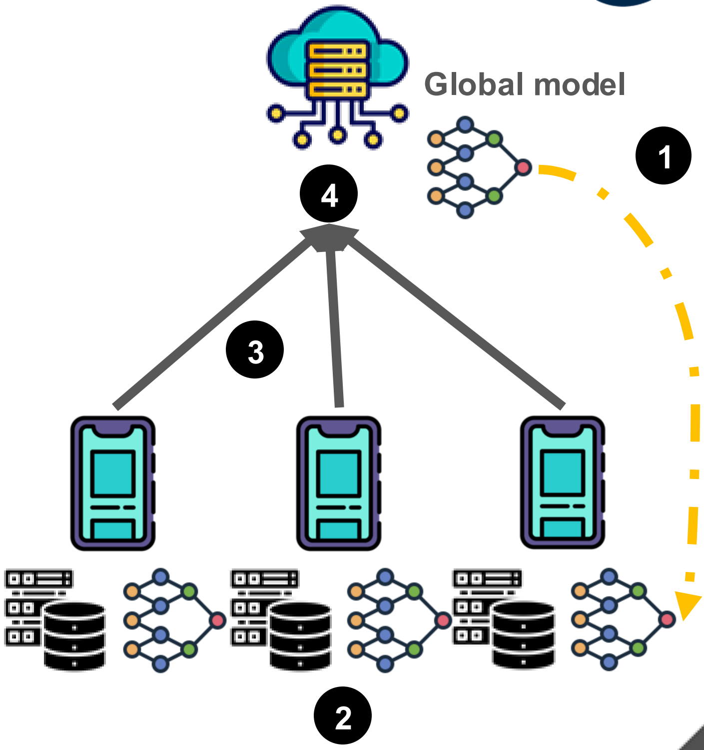

[MCM17] McMahan et al., **Communication-efficient learning of deep networks from decentralized data**. AISTATS'17.
[KAI21] Kairouz et al., **Advances and open problems in federated learning.** Fondations and Trends in Machine Learning'21.
[VAN24] Van Genderen et al., **Federated data access and federated learning: improved data sharing, AI model development, and learning in intensive care,** Intensive Care Medicine 2024.
[GU24] Gu et al., **Enhancing Data Provenance and Model Transparency in Federated Learning Systems--A Database Approach**, Preprint'24.

Gossip Learning [HEG19]

❶ **Stochastic Model Exchange:** Each user i sends model θ_i^t to its neighbors j ∈ $N(i)$

❷ **Local Aggregation and Training:** user i aggregates received models

$$\theta_i^{t+\frac{1}{2}} = \omega_{ii}\, \theta_i^t + \sum_{j \in N(i)} \omega_{ij} \theta_j^t$$

and updates locally

$$\theta_i^{t+1} = \theta_i^{t+\frac{1}{2}} - \eta \nabla L(\theta_i^{t+\frac{1}{2}}; D_i)$$

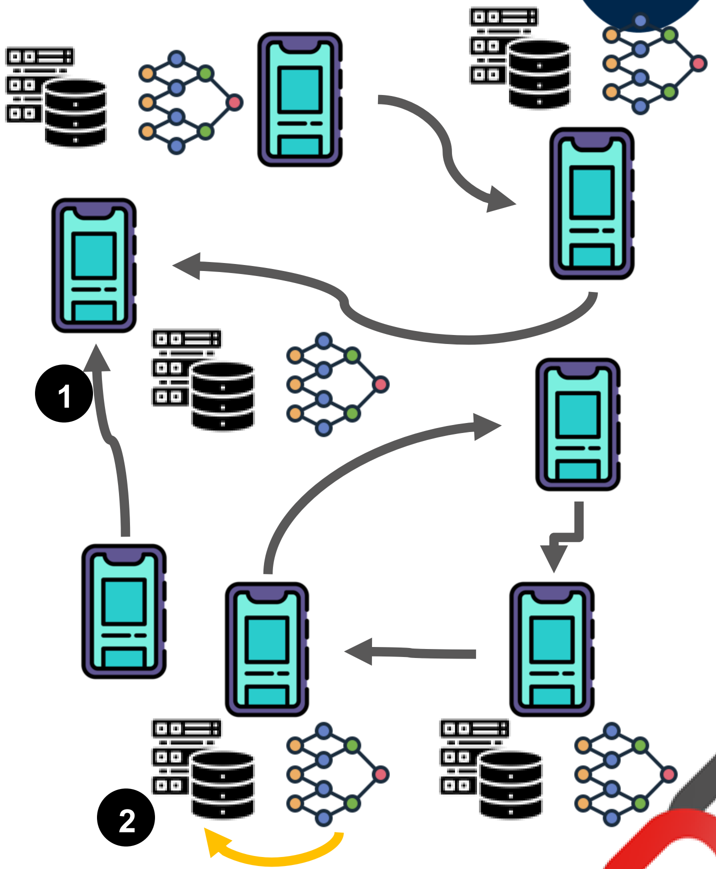

[HEG19] Hegedűs et al., **Gossip learning as a decentralized alternative to federated learning.** DAIS'19.

Gossip Learning [HEG19]

- **Graph dependence**

Consensus rate limited by graph topology [BOY06]

- **The need for dense graphs**

Faster convergence requires denser graphs

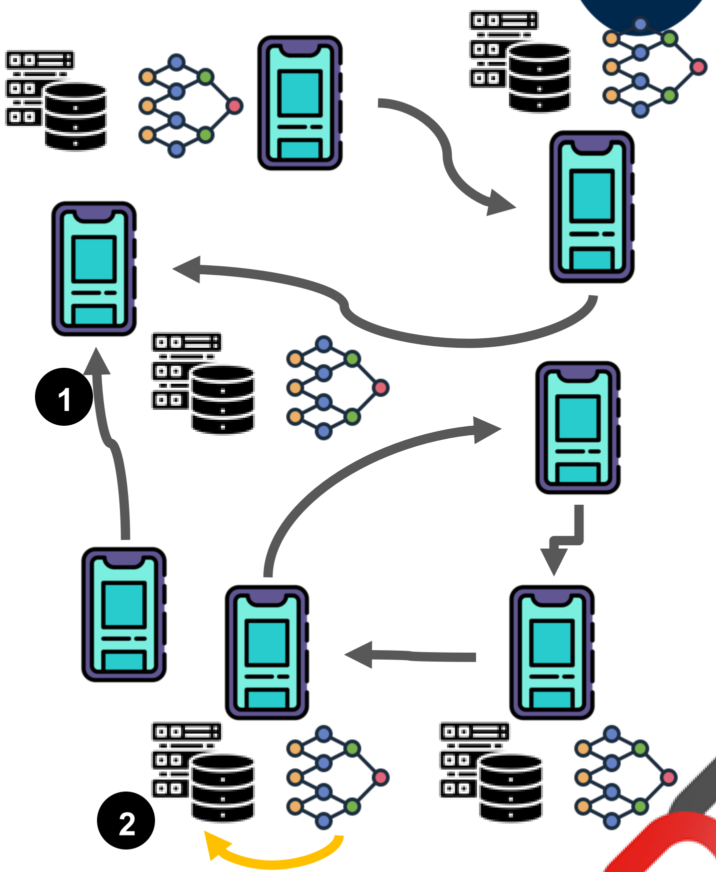

[HEG19] Hegedűs et al., **Gossip learning as a decentralized alternative to federated learning.** DAIS'19.
[BOY06] Boyd et al., **Randomized gossip algorithms.** IEEE Trans. Inf. Theory'06.

Dynamic Gossip Learning

3 Random Peer Sampling

Example Protocol: *View Shuffling* [BUS11]

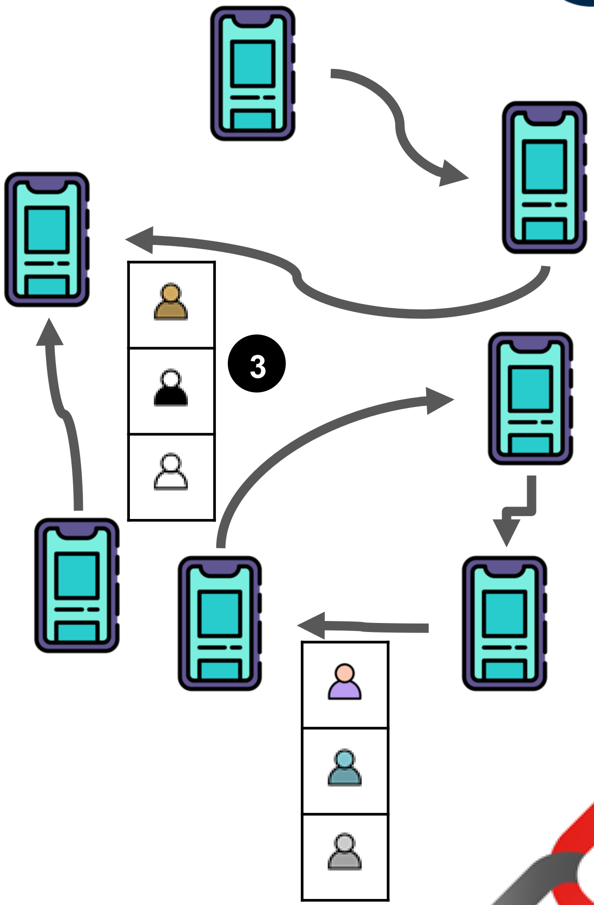

Properties

- Graph-size independent consensus rate [SON22]
- Exact-averaging with logarithmic degree graphs [YIN21]

[BUS11] Busnel et al., **On the uniformity of peer sampling based on view shuffling.** Journal of Parallel and Distributed Computing'11.
[SON22] Song et al., **Communication-efficient topologies for decentralized learning with o (1) consensus rate.** NeurIPS'22.
[YIN21] Ying et al., **Exponential graph is provably efficient for decentralized deep training.** NeurIPS'21.

Byzantine attacks

Open participation exposes the system to Byzantine users

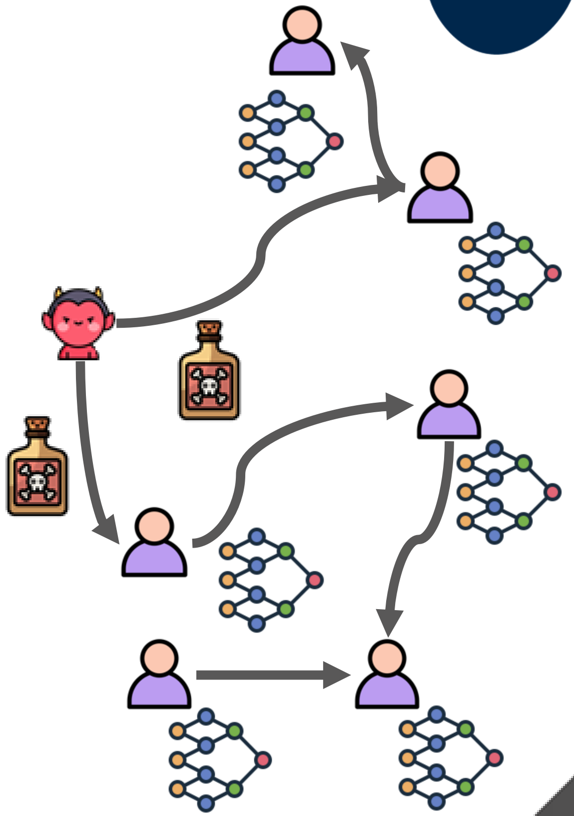

[GUE24] Guerraoui et al., **Byzantine machine learning: A primer.** ACM Computing Surveys'24.
[WAN20] Wang et al., **Attack of the tails: Yes, you really can backdoor federated learning.** NeurIPS'20.

Byzantine attacks

Open participation exposes the system to Byzantine users

- **Poisoning:** causes model divergence [GUE24]

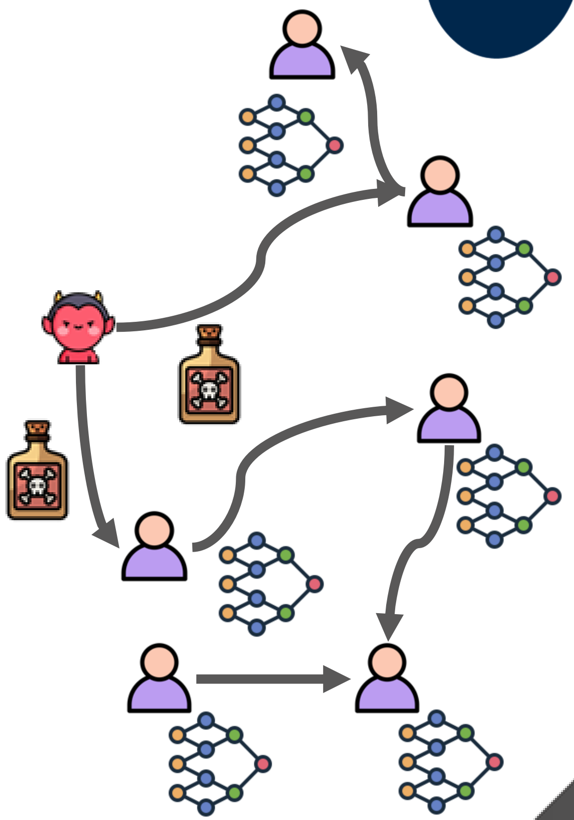

[GUE24] Guerraoui et al., **Byzantine machine learning: A primer.** ACM Computing Surveys'24.
[WAN20] Wang et al., **Attack of the tails: Yes, you really can backdoor federated learning.** NeurIPS'20.

Byzantine attacks

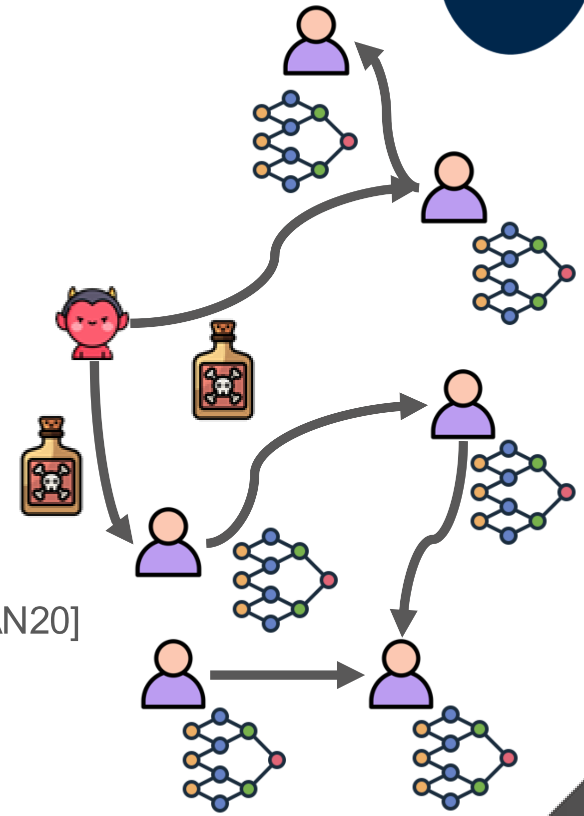

Open participation exposes the system to Byzantine users

- **Poisoning:** causes model divergence [GUE24]

- **Backdoor:** implants specific model misbehavior for [WAN20]

[GUE24] Guerraoui et al., **Byzantine machine learning: A primer.** ACM Computing Surveys'24.
[WAN20] Wang et al., **Attack of the tails: Yes, you really can backdoor federated learning.** NeurIPS'20.

Context: Poisoning attacks

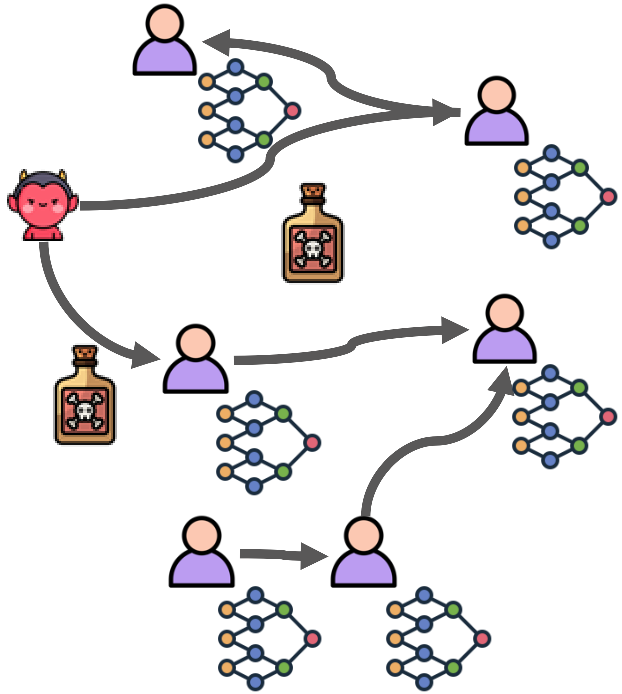

Context: Poisoning attacks

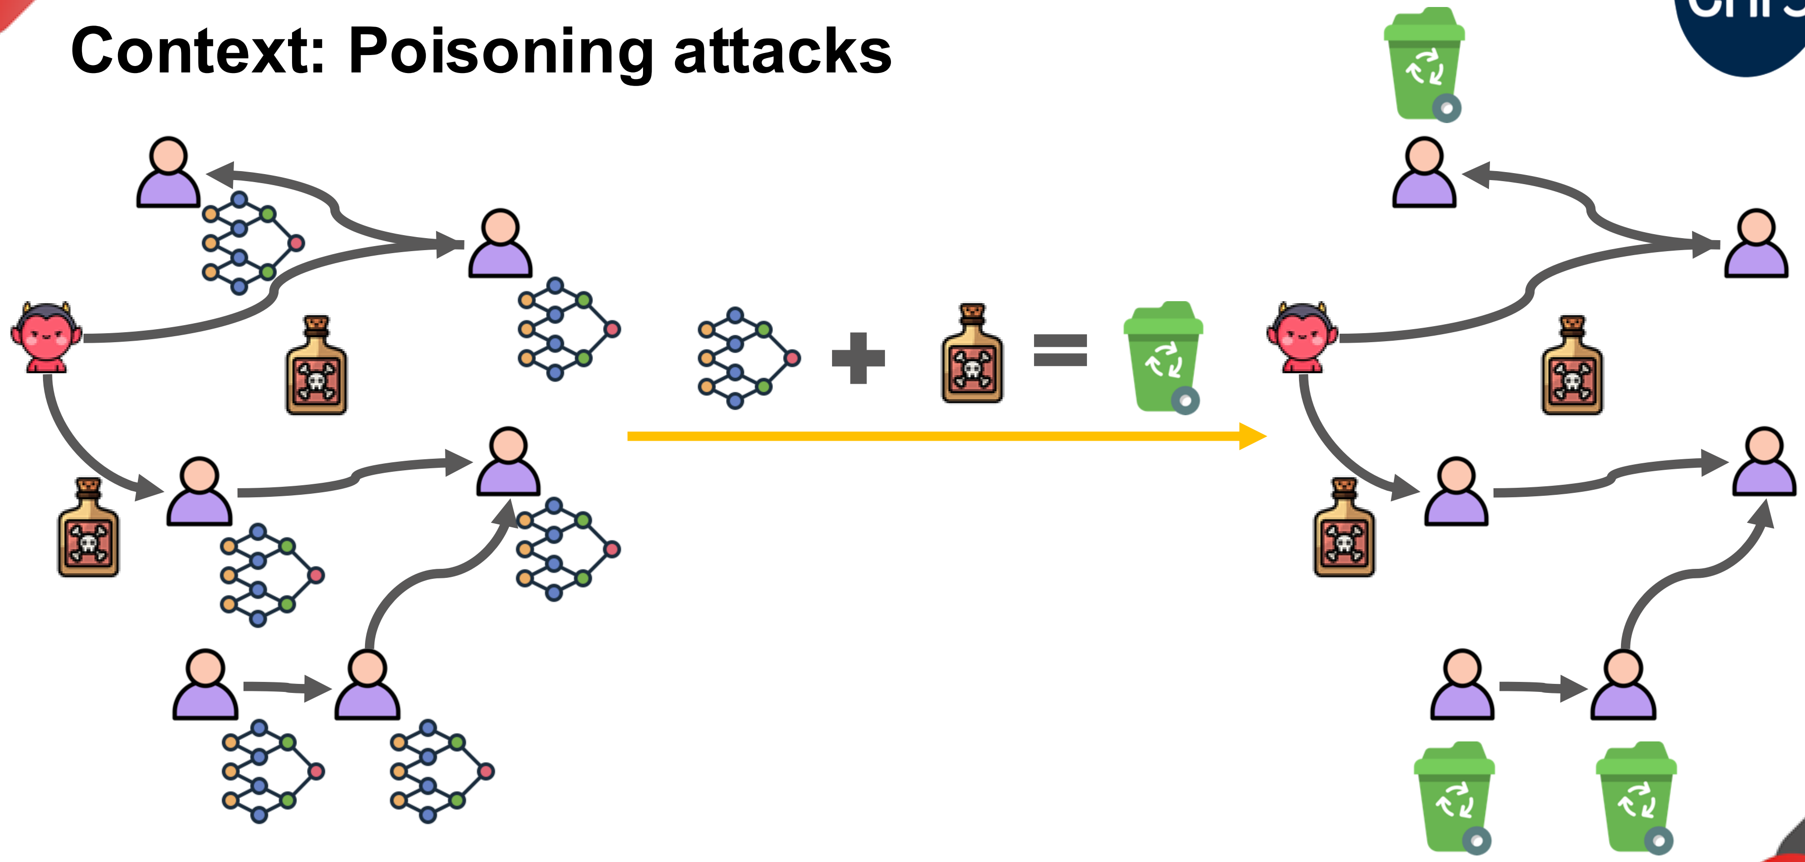

State of the Art: Poisoning defenses

- **Objective:** Filter or limit the impact of outlier models

- Vast literature in the federated setting [PIL22, ALL23]

Krum, Coordinate-wise trimmed median…

Not necessarily adapted to the Gossip Setting

- Rely on a large population of models
- Absence of considerations w.r.t the communication graph

[PIL22] Pillutla et al., **Robust aggregation for federated learning**. IEEE Trans. Sign. Proc.’22.
[ALL23] Allouah et al., **Fixing by mixing: A recipe for optimal byzantine ml under heterogeneity.** AISTATS ’23.

State of the Art: Robust aggregators in Gossip Learning

- **Same Objective:** Filter or limit the impact of outlier models

- **Key Properties:**
  - Consider the local model as a reference point
  - Consider the connectivity of the (honest) graph [FAN22]
  - Guarantees under some constraints (e.g., high connectivity)

- **Assumption:**
  - Known fixed threshold b: maximum number of byzantine nodes per neighbourhood [HE22, WU23]

[FAN22] Fang et al., **Bridge: Byzantine-resilient decentralized gradient descent.** IEEE Trans. Signal Inf. Process. Netw.'22.
[HE22] He et al., **Byzantine-robust decentralized learning via clippedgossip.** Preprint'22.
[WU23] Wu et al., **Byzantine-resilient decentralized stochastic optimization with robust aggregation rules.** IEEE Trans. Sign. Process.'23.

State of the Art: Robust aggregators

Clipped Summation [GAU25]:

For each neighbor $j \in N(i)^t$

1. Compu

2. Compu

3. Sort in … s

Details in the next presentation!

4. Aggregate $\theta_i^{t+1} = \theta_i^t + \sum_{k=1}^{v} \omega_k \cdot clip(z_r^t, \pi_i^t)$

where $\pi_i^t = ||z_{2b}^t||$ (the 2b-th largest norm)

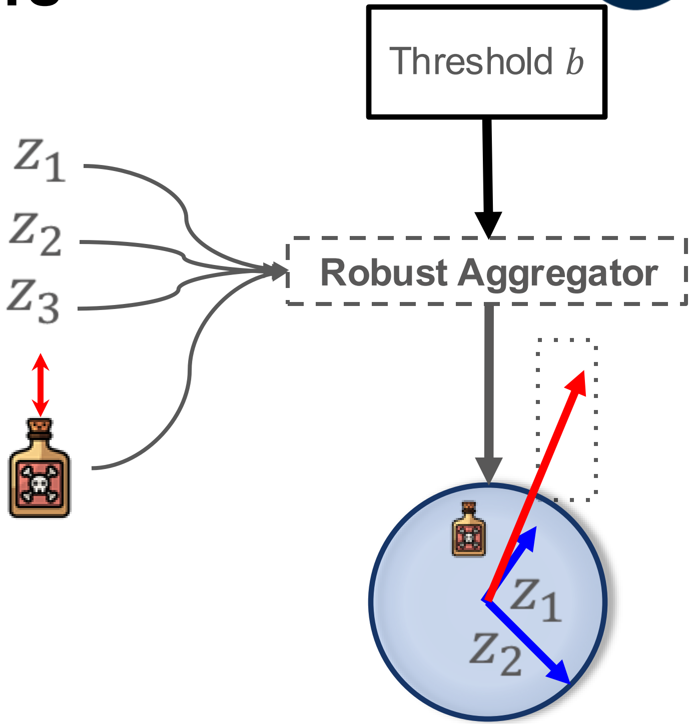

[GAU25] Gaucher et al., **Unified Breakdown Analysis for Byzantine Robust Gossip**. ICML'25

State of the Art: Limitation of Robust Aggregators

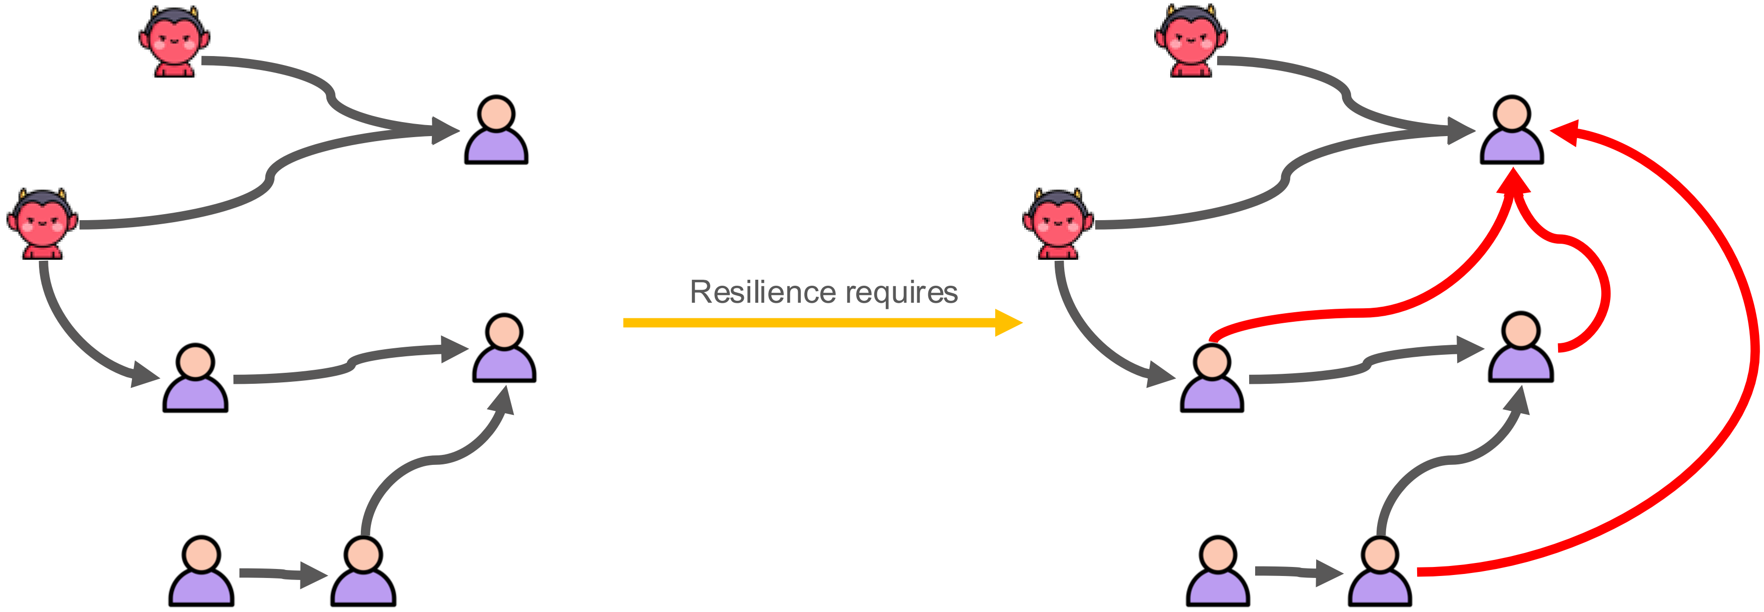

$b = 2, n = 7$

$\forall i \in N, |\mathrm{N(i)}| = 2b + 1 = 5$

Achieving worst-case resilience requires (extremely) dense graphs

State of the Art: Limitation of Robust Aggregators

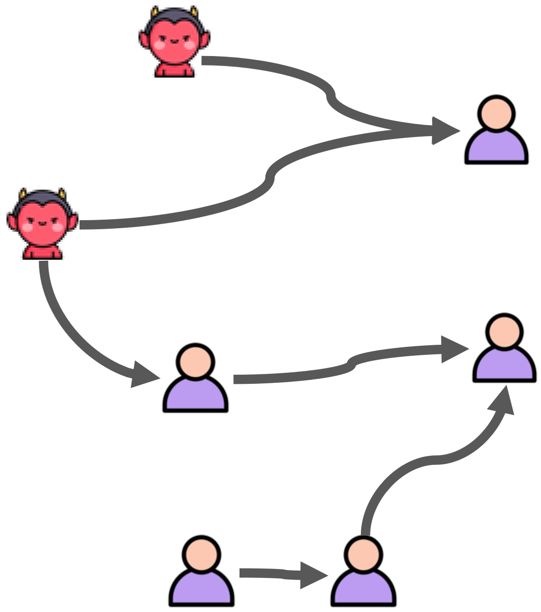

$b = 2, n = 7$

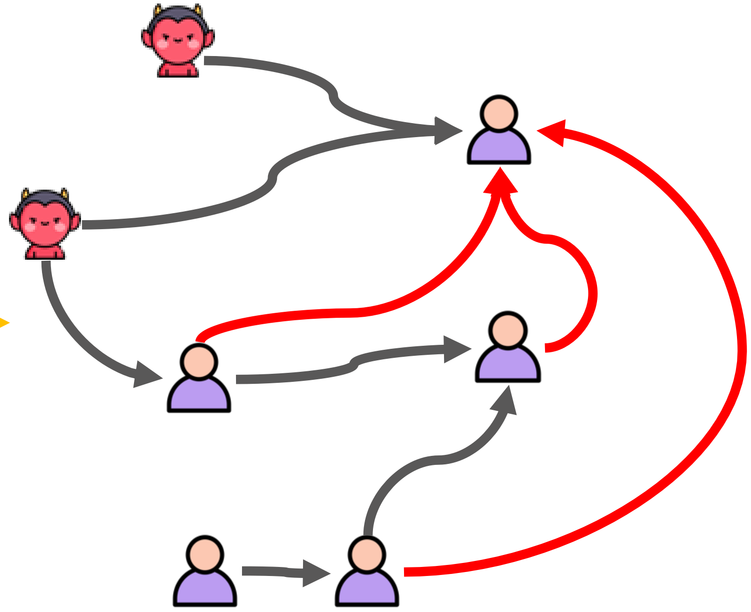

$\forall i \in N, |\mathrm{N}(i)| = 2b + 1 = 5$

Achieving worst-case resilience requires (extremely) dense graphs

Can Dynamic Gossip enable sparser graphs?

Context: Peer Sampling Flooding Attacks

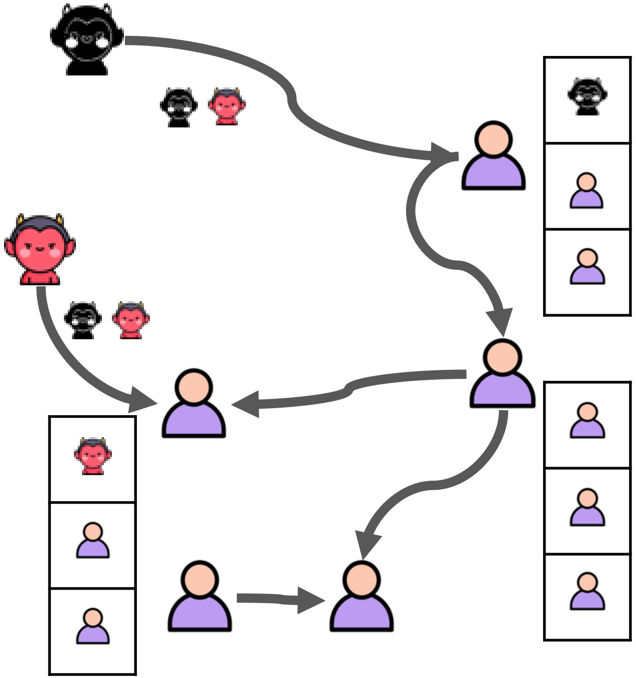

Context: Peer Sampling Flooding Attacks

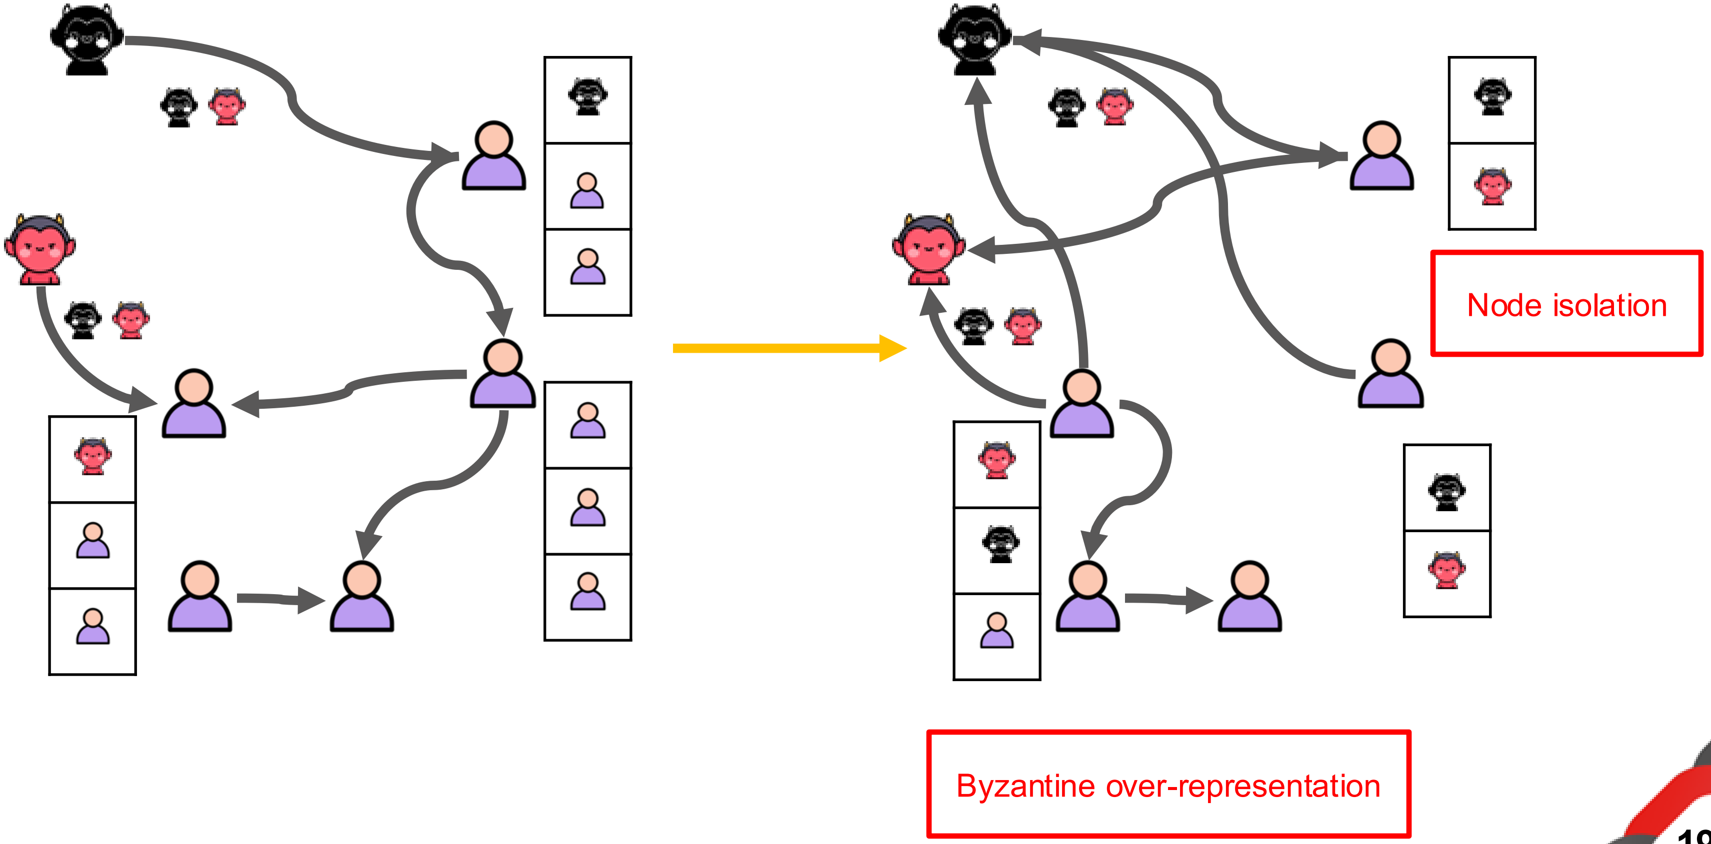

State of the Art: Byzantine-resilient Peer Sampling

- **Objective:** Peer discovery with resilience to attacks
- **Key Properties:**
  - Bound the probability of node isolation [BOR06, AUV23]
  - Ensure that the local Byzantine proportion tends toward the global one

Example: BASALT [AUV23]

- **Methodology:**
  - Peer identifiers are discovered through stochastic peer-to-peer exchanges
  - Local peer selection criterion based on uniform hash functions

[BOR08] Bortnikov et al., **Brahms: Byzantine resilient random membership sampling.** PODC'08.
[AUV23] Auvolat et al., **Basalt: A rock-solid byzantine-tolerant peer sampling for very large decentralized networks.** Middleware'23.

State of the Art: Byzantine-resilient Peer Sampling

- **Objective:** Peer discovery with resilience to attacks
- **Key Properties:**
  - Bound the probability of node isolation [BOR06, AUV23]
  - Ensure that the local Byzantine proportion tends toward the global one

Example: BASALT [AUV23]

  - **Methodology:**
    - Peer identifiers are discovered through stochastic peer-to-peer exchanges
    - Local peer selection criterion based on uniform hash functions
- **Applications:**
  - Message dissemination
  - File sharing, content discovery
  - Data replication

[BOR08] Bortnikov et al., **Brahms: Byzantine resilient random membership sampling.** PODC'08.
[AUV23] Auvolat et al., **Basalt: A rock-solid byzantine-tolerant peer sampling for very large decentralized networks.** Middleware'23.

How can Gossip Learning be made resilient to simultaneous Poisoning and Flooding attacks?

GRANITE (Big Picture)

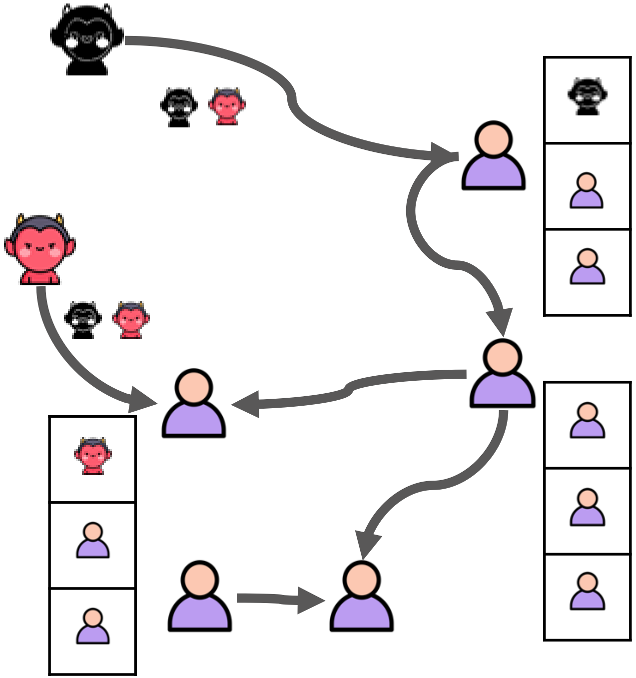

GRANITE (Big Picture)

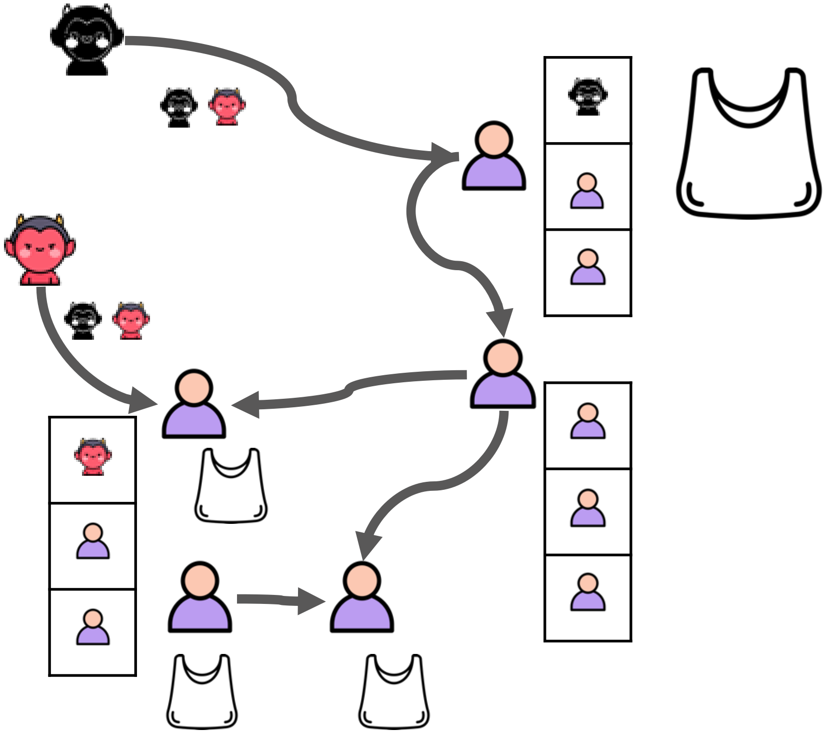

GRANITE (Big Picture)

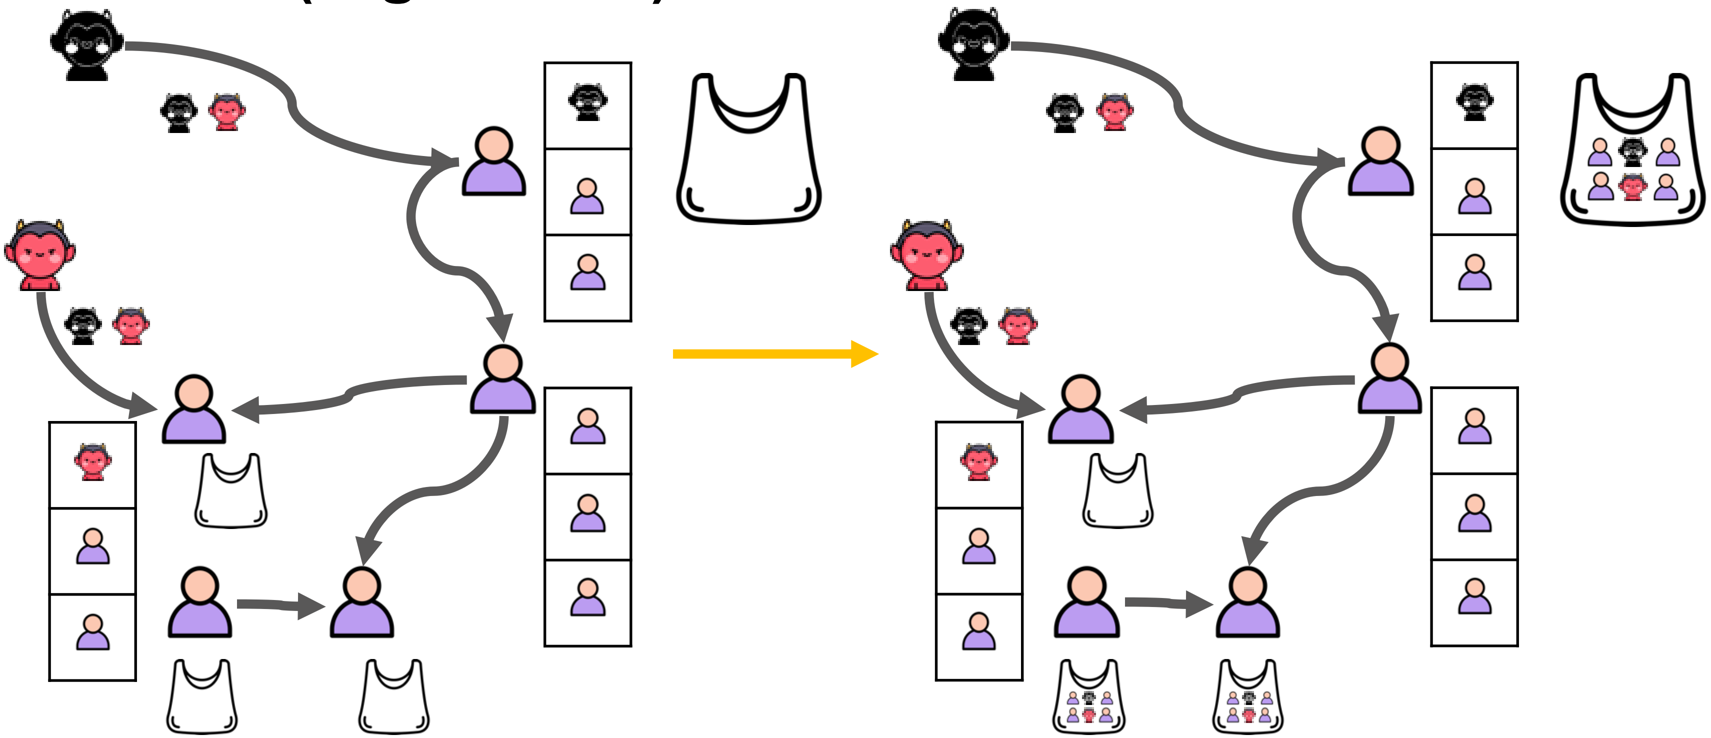

GRANITE (Big Picture)

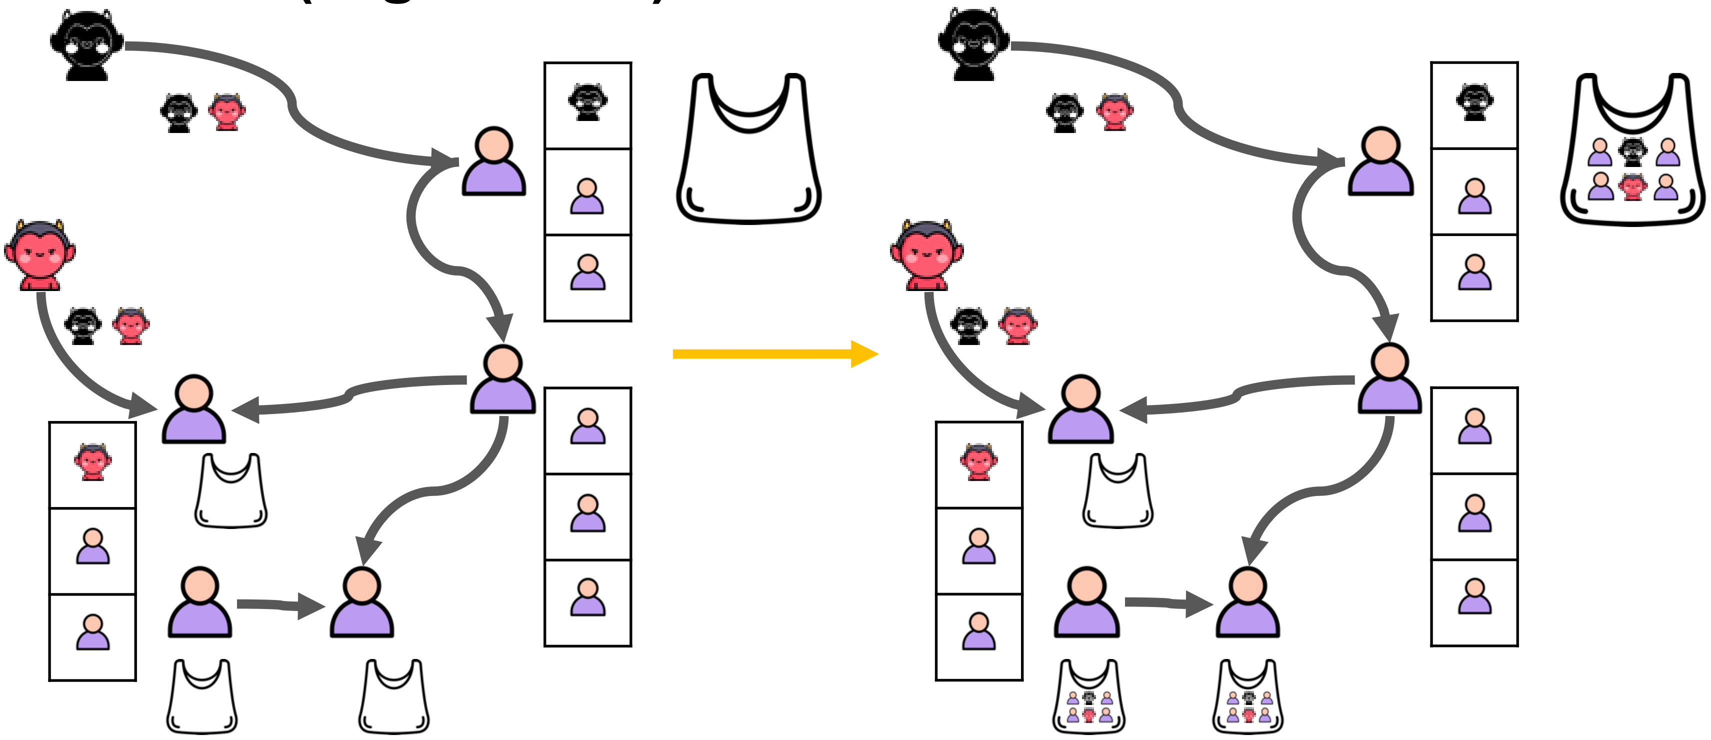

History-aware Peer Sampling

Adaptive Probabilistic Threshold

GRANITE: History-aware Peer Sampling

Estimates

$B(t)$
Local proportion of byzantine nodes

$\in$

Effective proportion of local byzantine nodes

Rate of arrival of honest identifiers

$$\mathbb{E}[b] \leq \mathrm{B}(t) = \frac{|\mathrm{B}|}{|\mathrm{B}| + |\mathrm{H}| - \left(|\mathrm{H}| - C(0)\right) \cdot e^{-\left(\frac{\alpha}{|\mathrm{H}|}\right)t}} \xrightarrow[t \to +\infty]{} f$$

Global number of byzantine nodes

Global number of honest nodes

Initial number of honest nodes

Guarantee: Bound the local Byzantine proportion using global parameters and system dynamic and ensure **exponential** decay

Granite: Adaptive Probabilistic Threshold

$B(t)$
Local proportion of byzantine nodes

$$b(t) = \min((1+\delta) \cdot v \cdot B(t), v - 1)$$

Safety margin

Neighborhood size

z_1

z_2

z_3

Robust Aggregator

$$P\left(X^t \geq (1+\delta) \cdot v \cdot B(t)\right) \leq \epsilon$$

Failure probability

Expected number of byzantine nodes

Guarantee: Robust aggregation with prob. $1 - \epsilon$

GRANITE: Experiments

Experiments aim at answering the following questions:

- How resilient is GRANITE against combined Poisoning and Flooding Attacks?
- How does GRANITE compare to SotA Byzantine-resilient Peer Sampling protocols?
  - *Competitor:* BASALT [AUV23]

[AUV23] Auvolat et al., **Basalt: A rock-solid byzantine-tolerant peer sampling for very large decentralized networks.** Middleware'23.

GRANITE: Experimental Setting

- **Datasets:**
  - Purchase-100, MNIST (Heterogeneity with Dirichlet method $\beta = .5$)
- **Models:**
  - fully connected models, convolution network
- **Robust aggregator:** Clipped Summation
- **Poisoning Attack:** Fall of Empires [XIE21]
- Flooding attack
- Byzantine fractions of 0.1 and 0.3
- **Metrics:**
  - F1-Score
  - Honest Subgraph Strongly Connected Component Ratio (HSSR)

[XIE21] Xie et al., **Fall of empires: Breaking byzantine-tolerant sgd by inner product manipulation.** UAI'21.

GRANITE versus BASALT

- Dataset: Purchase-100
- 300 users
- 10% byzantine nodes

- Three CS parameterization under BASALT:
  - **Conservative:** $b = v - 1$
  - **Medium:** $b = 6 \cdot f \cdot v$
  - **Loose:** $b = 4 \cdot f \cdot v$

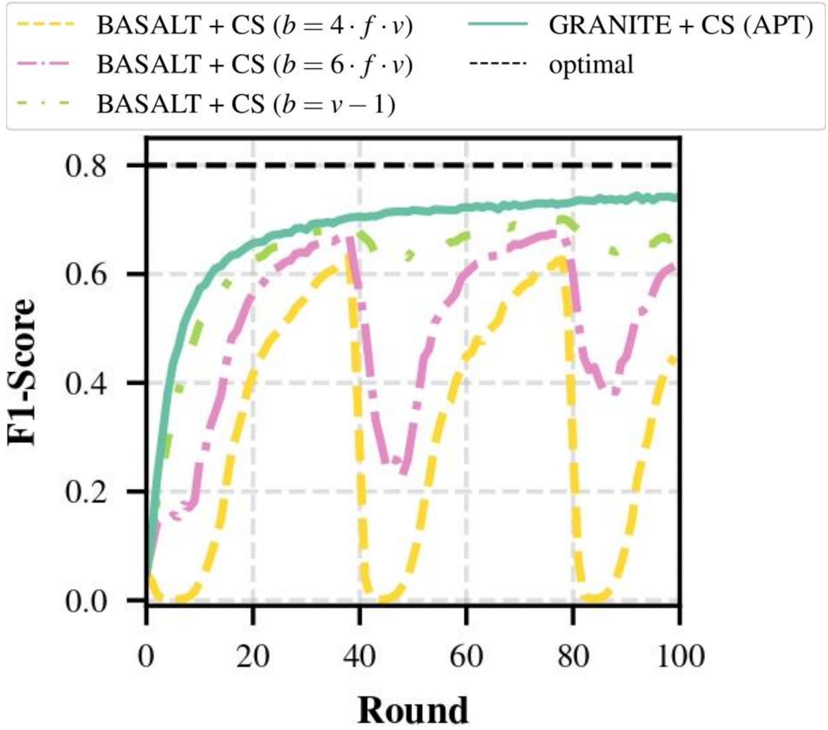

GRANITE versus BASALT

- Dataset: Purchase-100
- 300 users
- 10% byzantine nodes

- Three CS parameterization under BASALT:
  - **Conservative:** $b = v - 1$
  - **Medium:** $b = 6 \cdot f \cdot v$
  - **Loose:** $b = 4 \cdot f \cdot v$

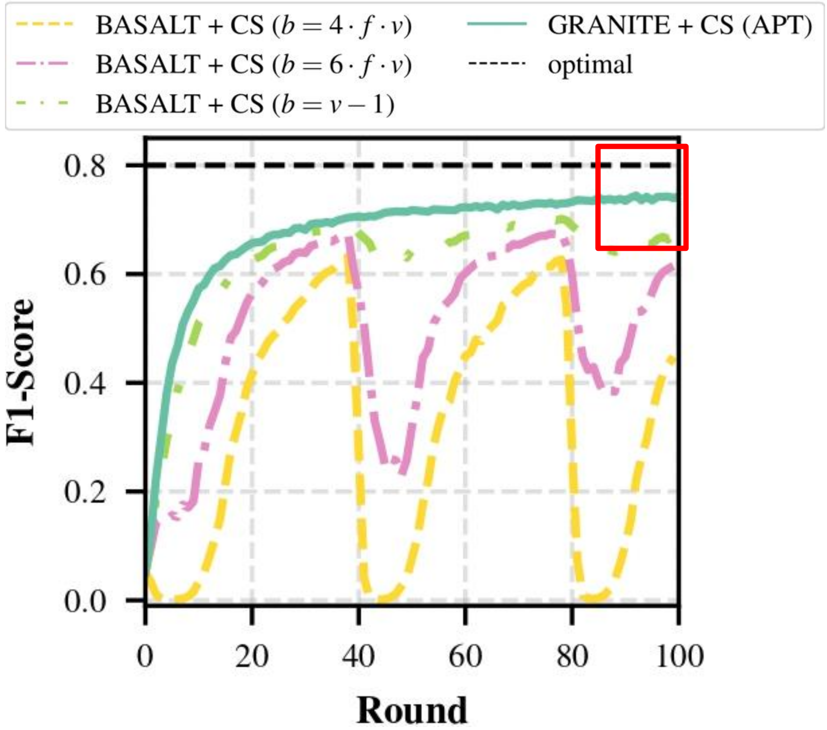

GRANITE converges in a stable fashion

GRANITE versus BASALT

- Dataset: Purchase-100
- 300 users
- 10% byzantine nodes

- Three CS parameterization under BASALT:
  - **Conservative:** $b = v - 1$
  - **Medium:** $b = 6 \cdot f \cdot v$
  - **Loose:** $b = 4 \cdot f \cdot v$

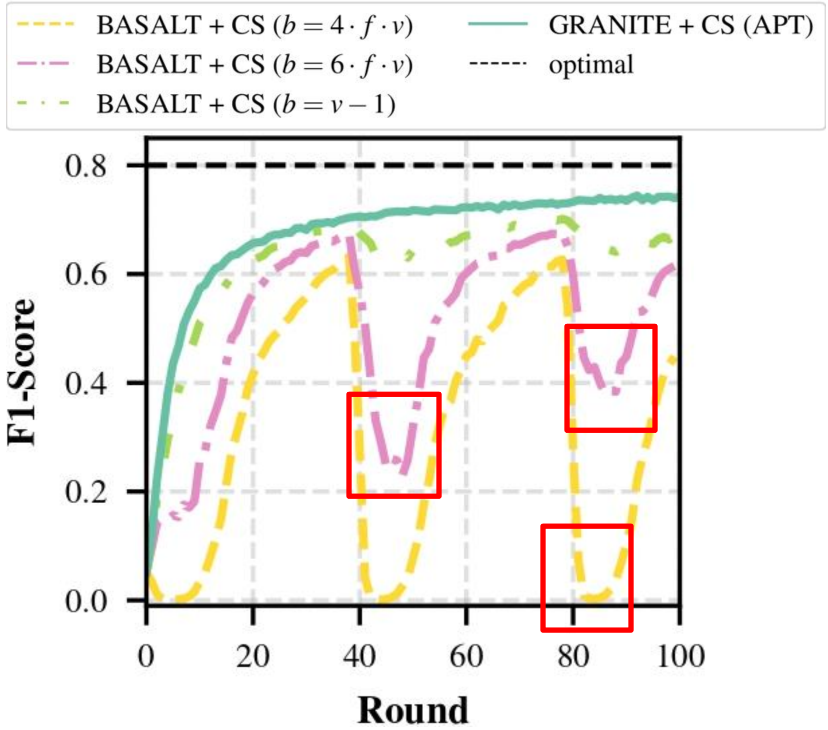

BASALT suffers major fluctuations and periodical valleys

GRANITE versus BASALT

- Dataset: MNIST
- 300 users
- 30% byzantine nodes

- Three CS parameterization under BASALT:
  - **Conservative:** $b = v - 1$
  - **Medium:** $b = 6 \cdot f \cdot v$
  - **Loose:** $b = 4 \cdot f \cdot v$

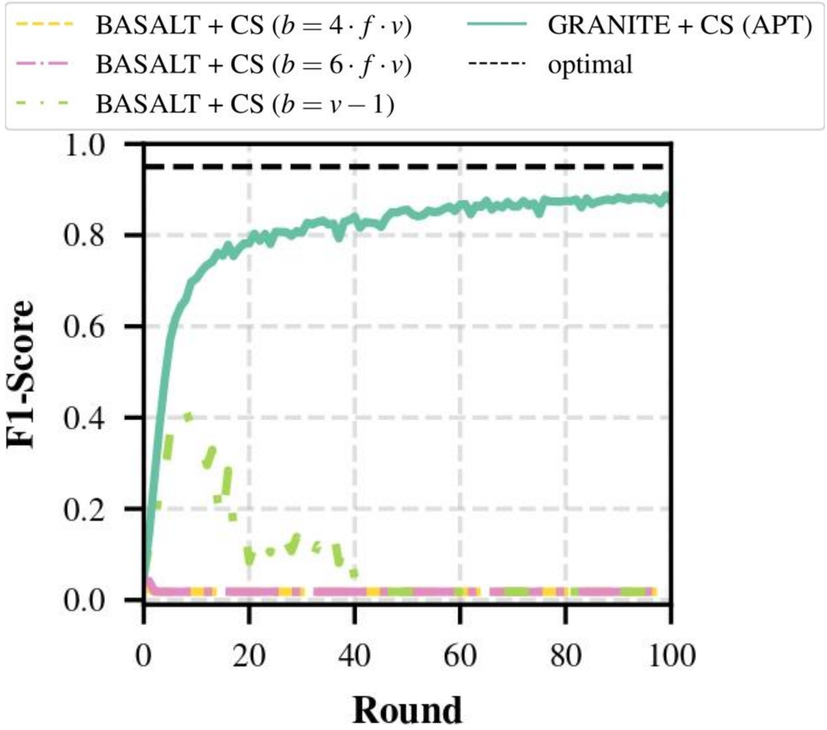

GRANITE versus BASALT

- Dataset: MNIST
- 300 users
- 30% byzantine nodes

- Three CS parameterization under BASALT:
  - **Conservative:** $b = v - 1$
  - **Medium:** $b = 6 \cdot f \cdot v$
  - **Loose:** $b = 4 \cdot f \cdot v$

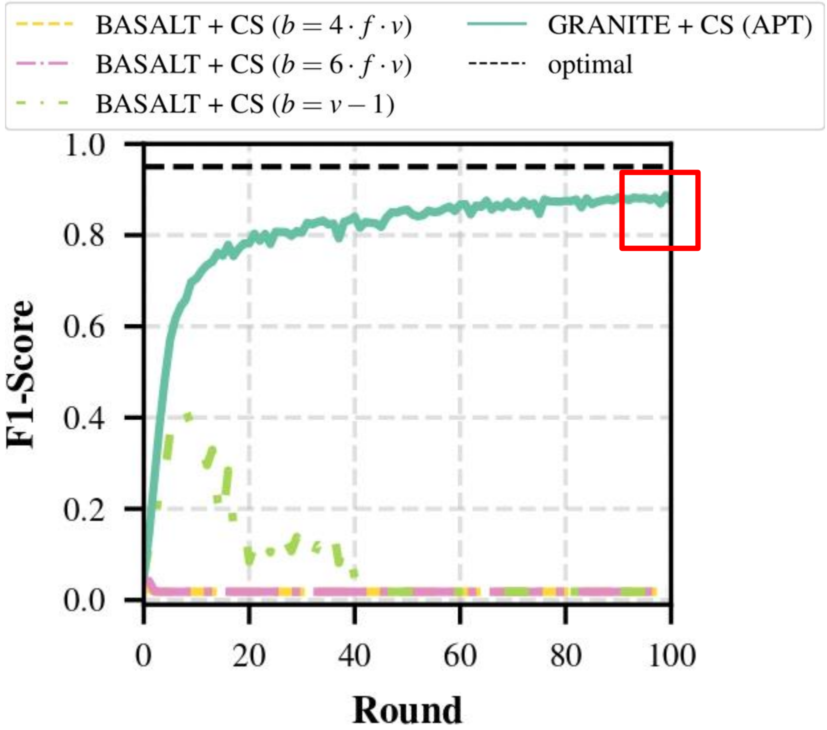

GRANITE converges towards the optimal performance

GRANITE versus BASALT

- Dataset: MNIST
- 300 users
- 30% byzantine nodes

- Three CS parameterization under BASALT:
  - **Conservative:** $b = v - 1$
  - **Medium:** $b = 6 \cdot f \cdot v$
  - **Loose:** $b = 4 \cdot f \cdot v$

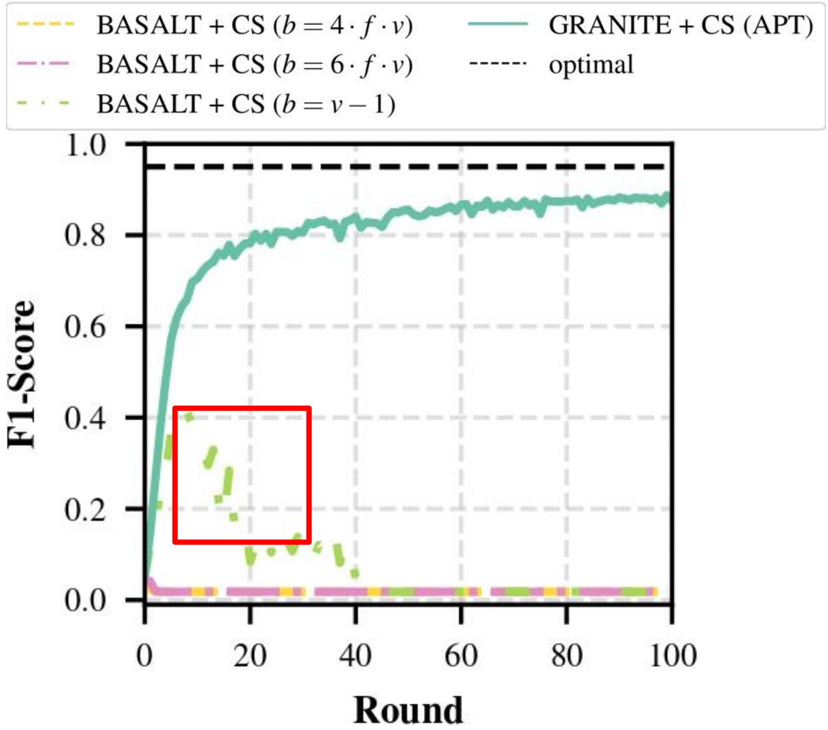

BASALT starts diverging as early as the 10th round

GRANITE: Conclusion

- Robust aggregators often require dense graphs
- Byzantine-resilient peer sampling have a different design context
- GRANITE bridges the gap between Byzantine-resilient peer sampling protocols and robust aggregators

Y. Belal, M. Maouche, S. Ben Mokhtar, & A. Simonet-Boulogne.
GRANITE: a Byzantine-Resilient Dynamic Gossip Learning Framework.
Preprint: https://arxiv.org/pdf/2504.17471

</>

GRANITE: a Byzantine-resilient Gossip Learning Framework.
https://anonymous.4open.science/r/Granite-Byzantine-Resilient-Dynamic-Gossip-Learning-Framework-4886