SONAR: Benchmarking Topology in Decentralized Learning

Exposing network structure as a first-class systems variable

MLSys 2026

Abhishek Singh*, Joyce Yuan*, Yichuan Shi* (MIT)
Rishi Sharma, Jonas Blanc, Martin Jaggi (EPFL)
Ramesh Raskar (MIT)

The Problem: From Centralized to Decentralized

- → **Traditional FL:** Star topology with central server = orchestration bottleneck, single point of failure, limited to one administrative domain
- → **Shift to P2P:** Clients exchange models directly, removing bottleneck but creating complex systems problem
- → **Communication topology** becomes the defining systems variable (rings, grids, random graphs, adaptive)
- → **Confounded factors:** Topology, data distribution, local algorithms, # clients, and attackers are coupled—must be studied together
- → **The Gap:** Existing frameworks (FedML, FedScale, Flower, DecentralizePy) benchmark accuracy but offer limited control over topology
- → No infrastructure to systematically analyze how these coupled variables interact

How can we **systematically study** the impact of **topology**, coordination, and system constraints on decentralized learning?

SONAR enables **systematic, reproducible** evaluation of these coupled interactions

SONAR: Topology as a First-Class Variable

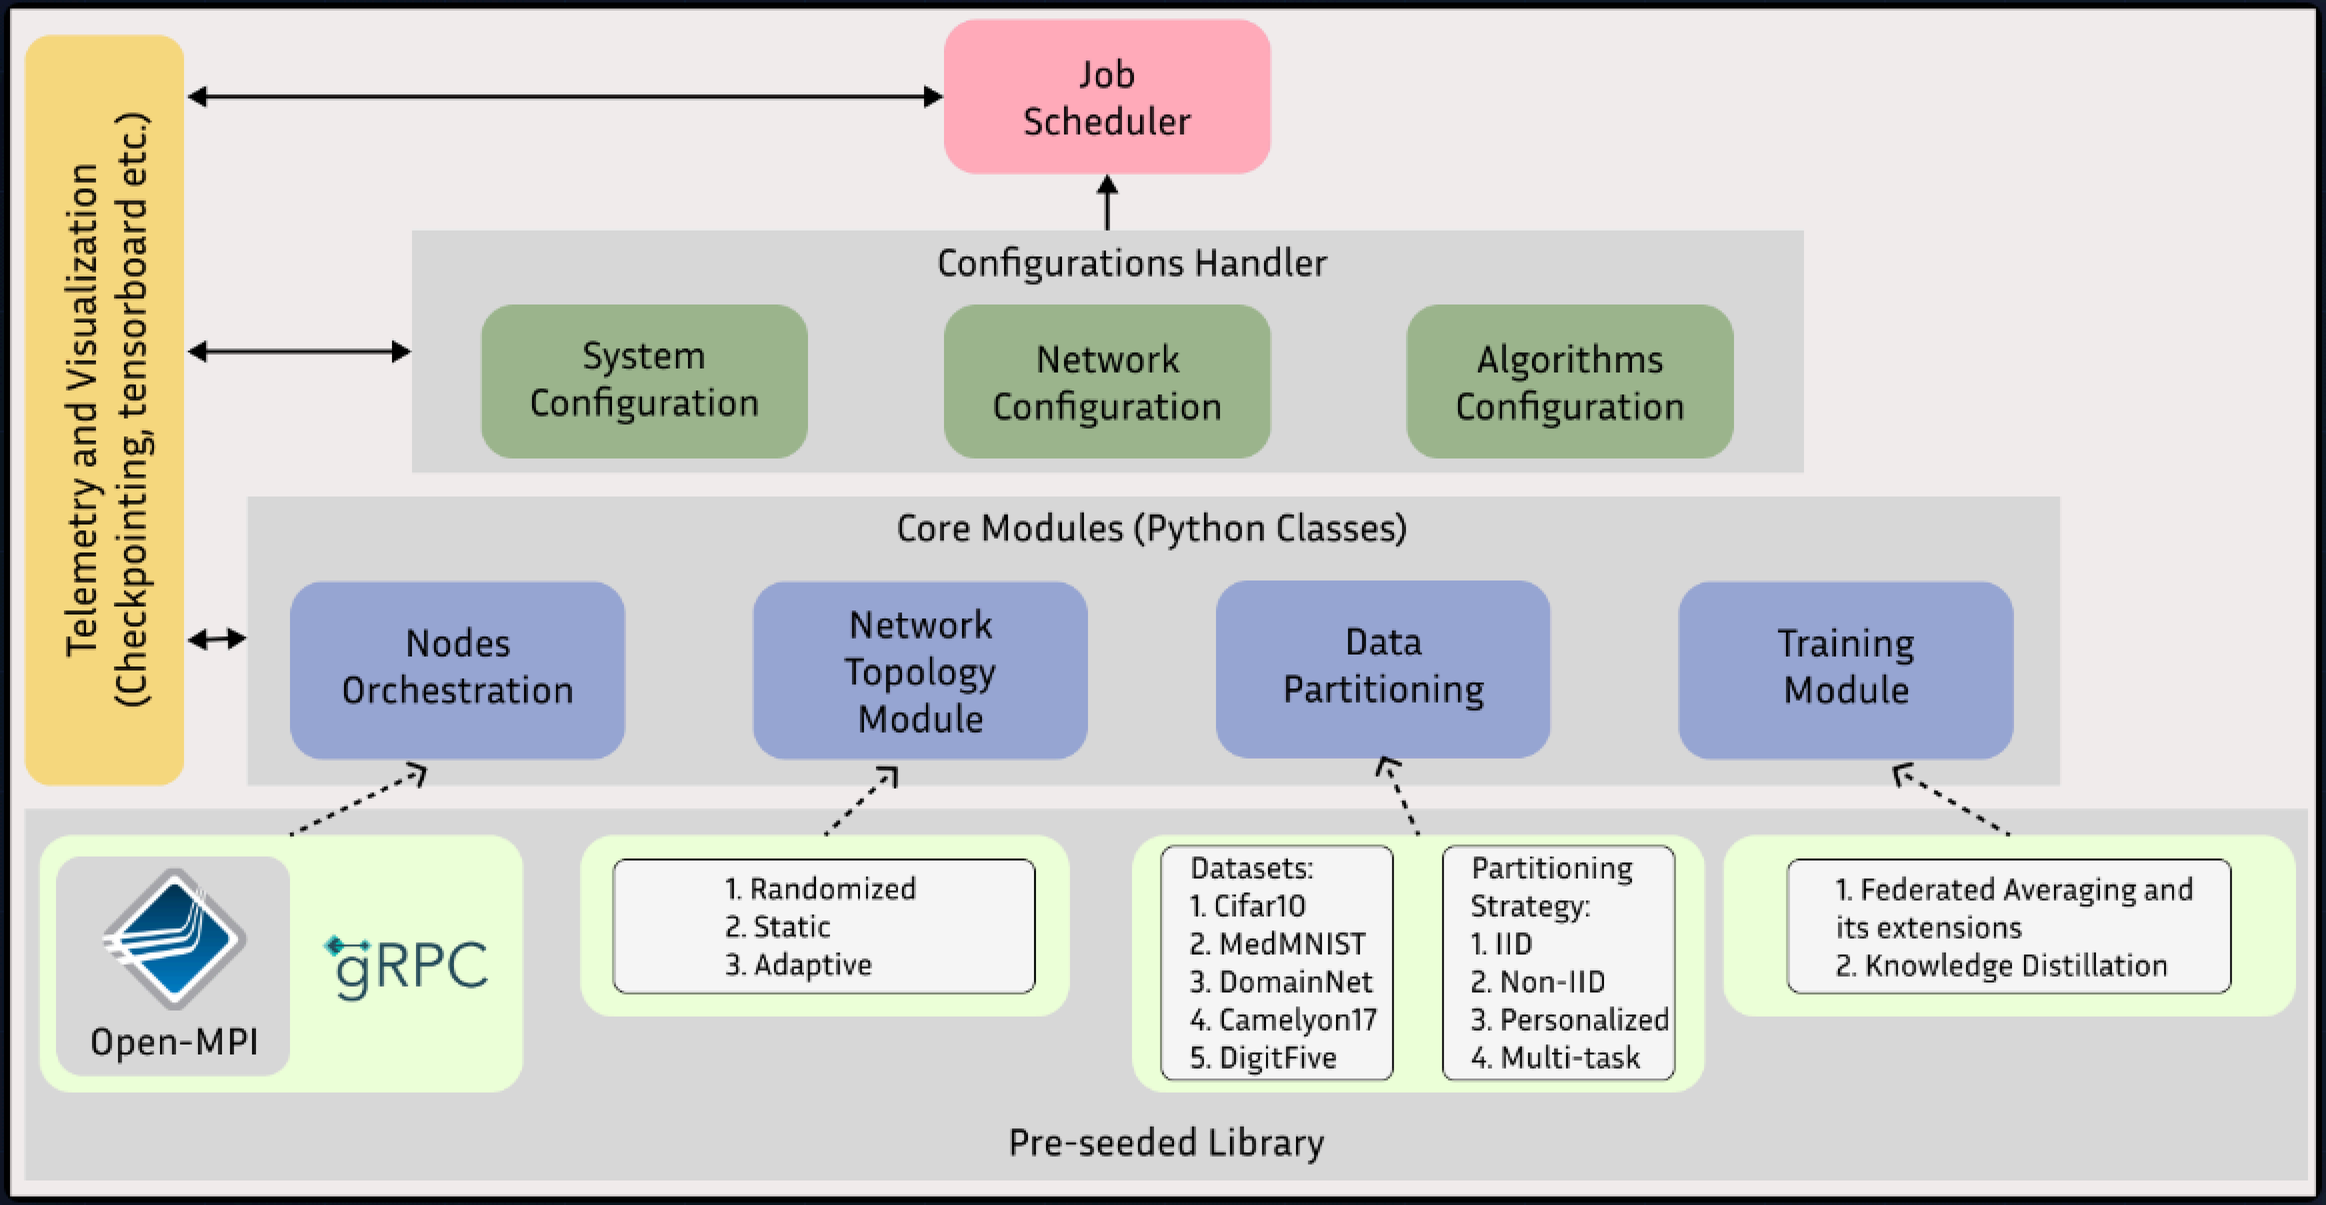

Four-layer architecture: orchestration, topology engine, communication, telemetry

→ Universal API, multiple backends (gRPC, MPI, WebRTC)

→ **Key differentiator:** Exposes topology as controllable, measurable dimension alongside data distribution and algorithms

Design: Universal API & Configuration

- **Modular architecture** with pluggable backends (gRPC, MPI, WebRTC)
- **Declarative configuration** for reproducibility across experiments
- **Universal API** abstracts communication complexity from algorithms
- **Built-in telemetry** for comprehensive systems performance analysis

Sample Configuration

```
sample_config: ConfigType = {
  "exp_id": "test_experiment",
  "num_users": 10,
  "session_id": 1111,
  "num_collaborators": ALL,
  "comm": {"type": "gRPC"},
  "dset": CIFAR10_DSET,
  "train_label_distribution": "iid",
  "topo": "ring",
  ...
}
```

Universal Communication API

```
send_update()
# transmit local weights or statistics

receive_updates()
# collect neighbor updates

aggregate()
# combine and apply received updates
```

Experimental Setup

Datasets

DomainNet (6 domains)

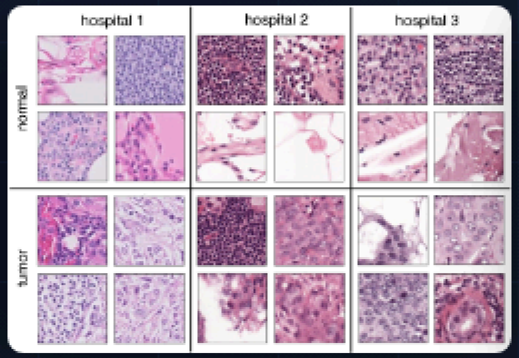

Camelyon17 (5 hospitals)

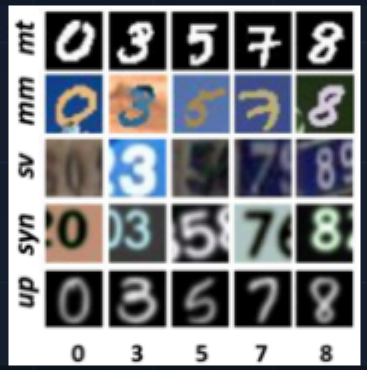

Digit-Five (5 sources)

Metrics

→ Test accuracy & AUC

→ Communication volume (GB/node)

→ Robustness to attacks

→ Privacy leakage (MIA, GIA)

Network Topologies

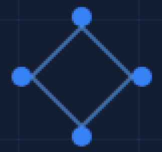

Ring

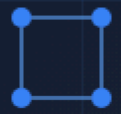

Torus

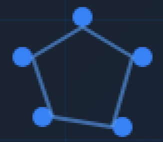

Random (ER)

Domain-aligned

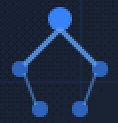

Adaptive (Top-K)

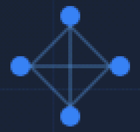

Fully Connected (FL)

Finding 1: Topology Choice Amplifies with Scale and Heterogeneity

→ **Scale matters:** Topologies start similar (36 users) but diverge dramatically by 225 users—gap grows from ~0.3 to ~0.9 AUC

→ **Data heterogeneity amplifies:** Domain-aligned topology advantage grows from ~9 AUC (2 domains) to ~13 AUC (6 domains)

→ **Topology is a first-order variable:** Determines both scalability and how effectively heterogeneous knowledge propagates

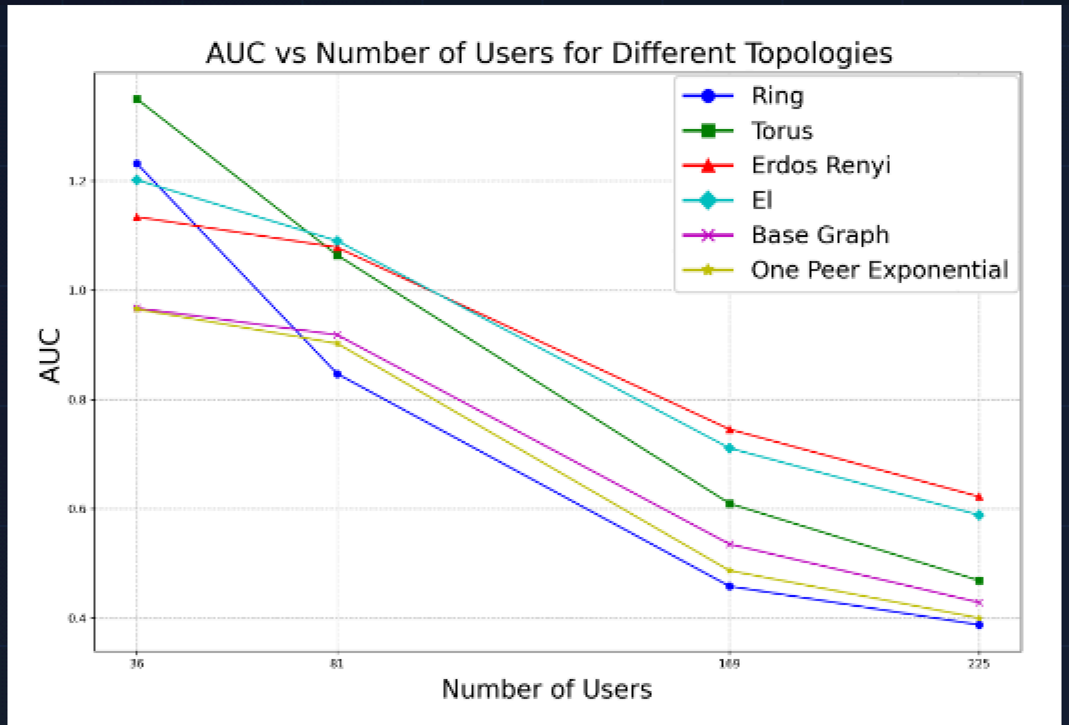

Topology divergence at scale: performance gap widens from 36 to 225 users

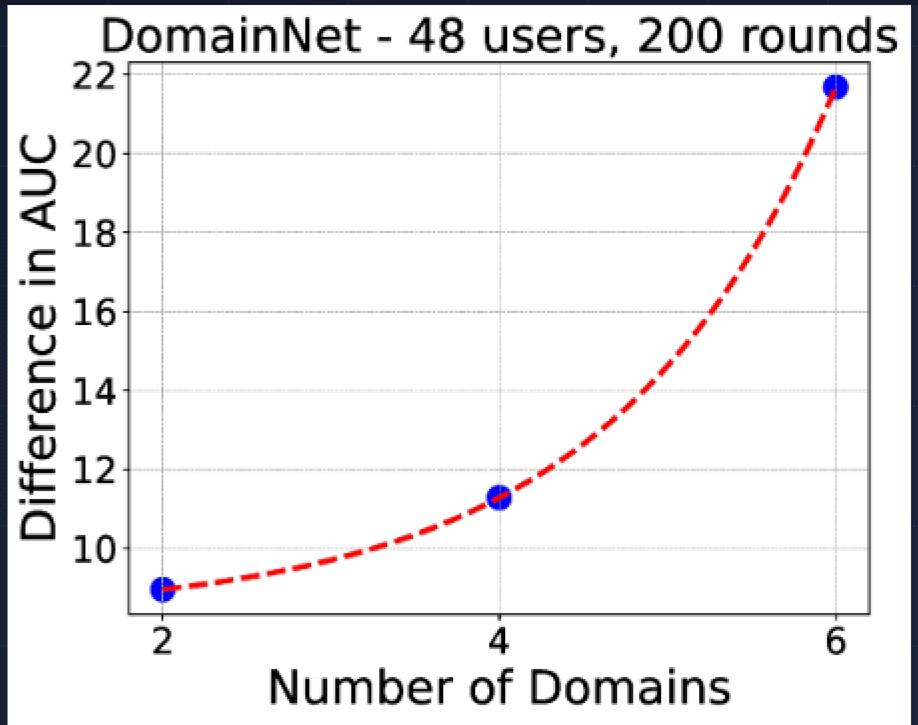

Heterogeneity amplification: advantage grows with more domains

Finding 2: Communication-Performance Tradeoffs

- **Federated Learning baseline:** Fully connected topology (64.40 GB/node/round) achieves highest accuracy but at massive communication cost
- **Structured sparse topologies win:** Ring (3.66 GB), Torus (7.36 GB), EI (11.1 GB) achieve 85-95% of FL accuracy at 5.7-17% of communication cost
- **Key insight:** More edges ≠ better convergence; sparse structure + data alignment = efficiency

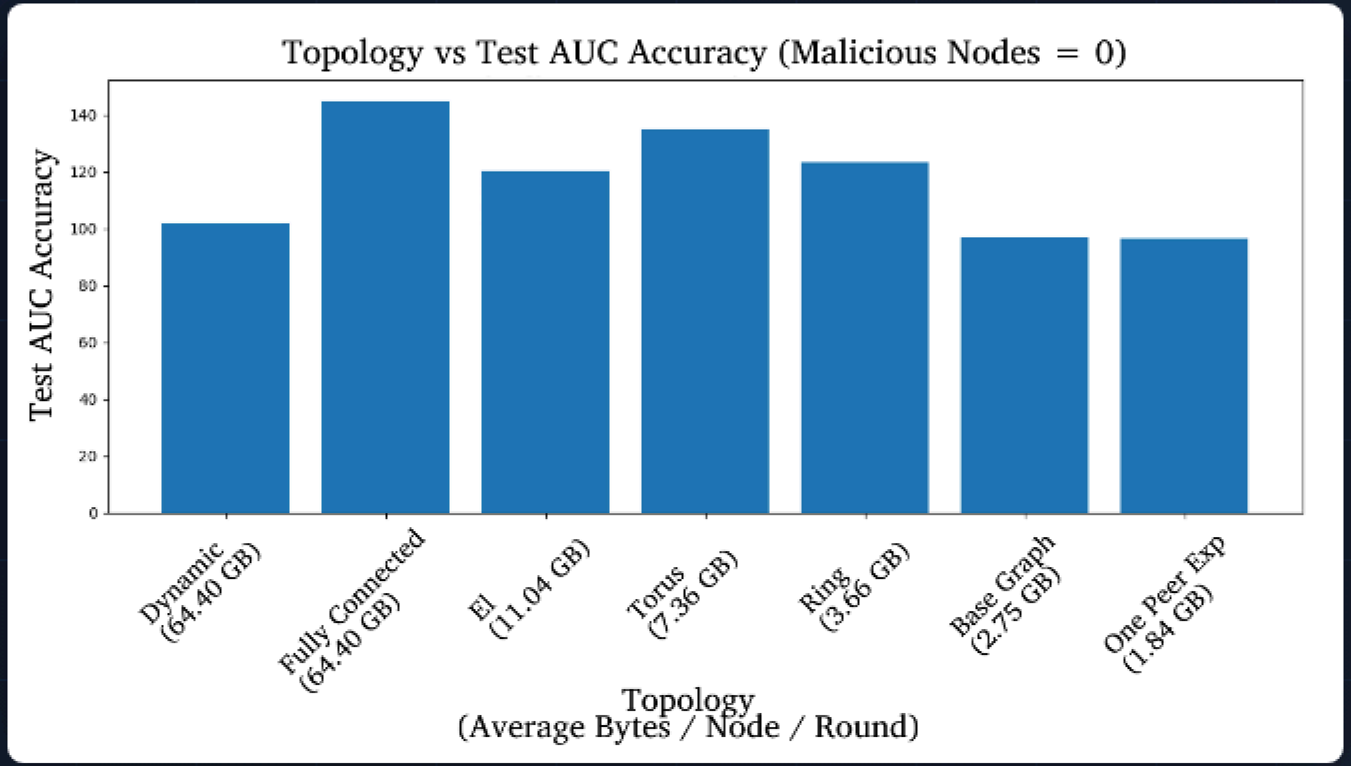

Efficiency frontier: structured topologies reach comparable AUC with dramatically lower communication volume

Finding 3: Adaptive Selection Requires Balancing Similarity and Diversity

→ **Intuition:** Select collaborators based on similarity metrics (gradient alignment, loss correlation) to improve accuracy

→ **Problem:** Greedy selection creates feedback loop → users repeatedly pick same collaborators → clusters collapse

→ **Top-K tradeoff:** Too small K (1-3) → collapse; optimal K ≈ domain size (12); too large K (18+) → harmful cross-domain mixing

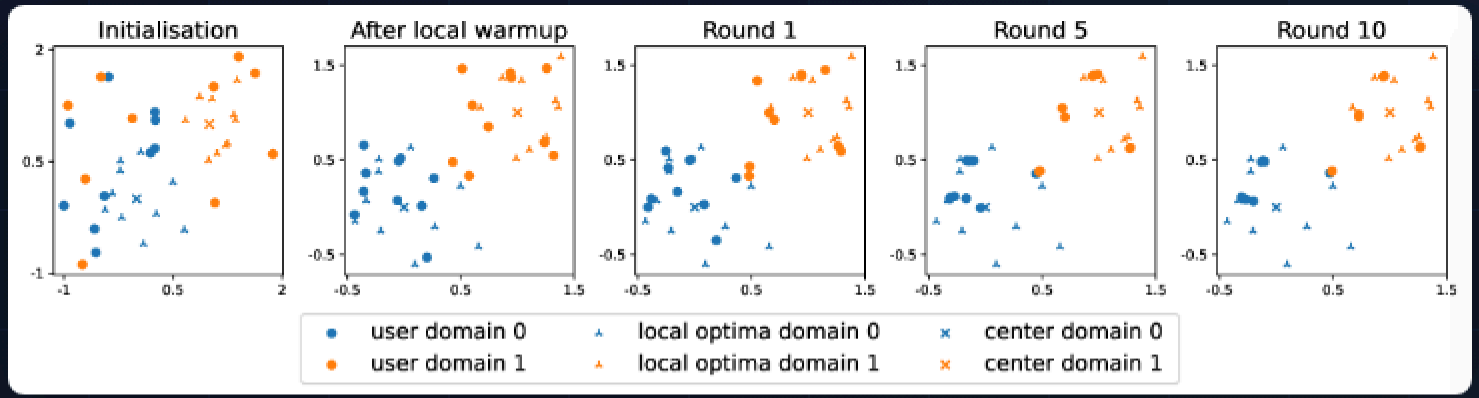

Simulation: users in 2D space converge into small, isolated clusters over training rounds

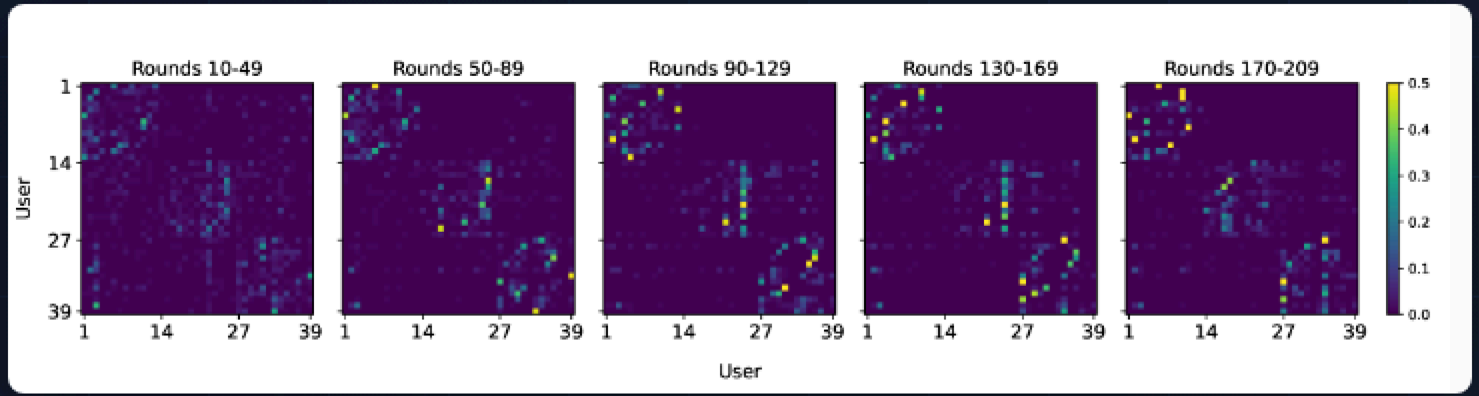

Real training: collaborator weights concentrate on few partners (bright spots) over time

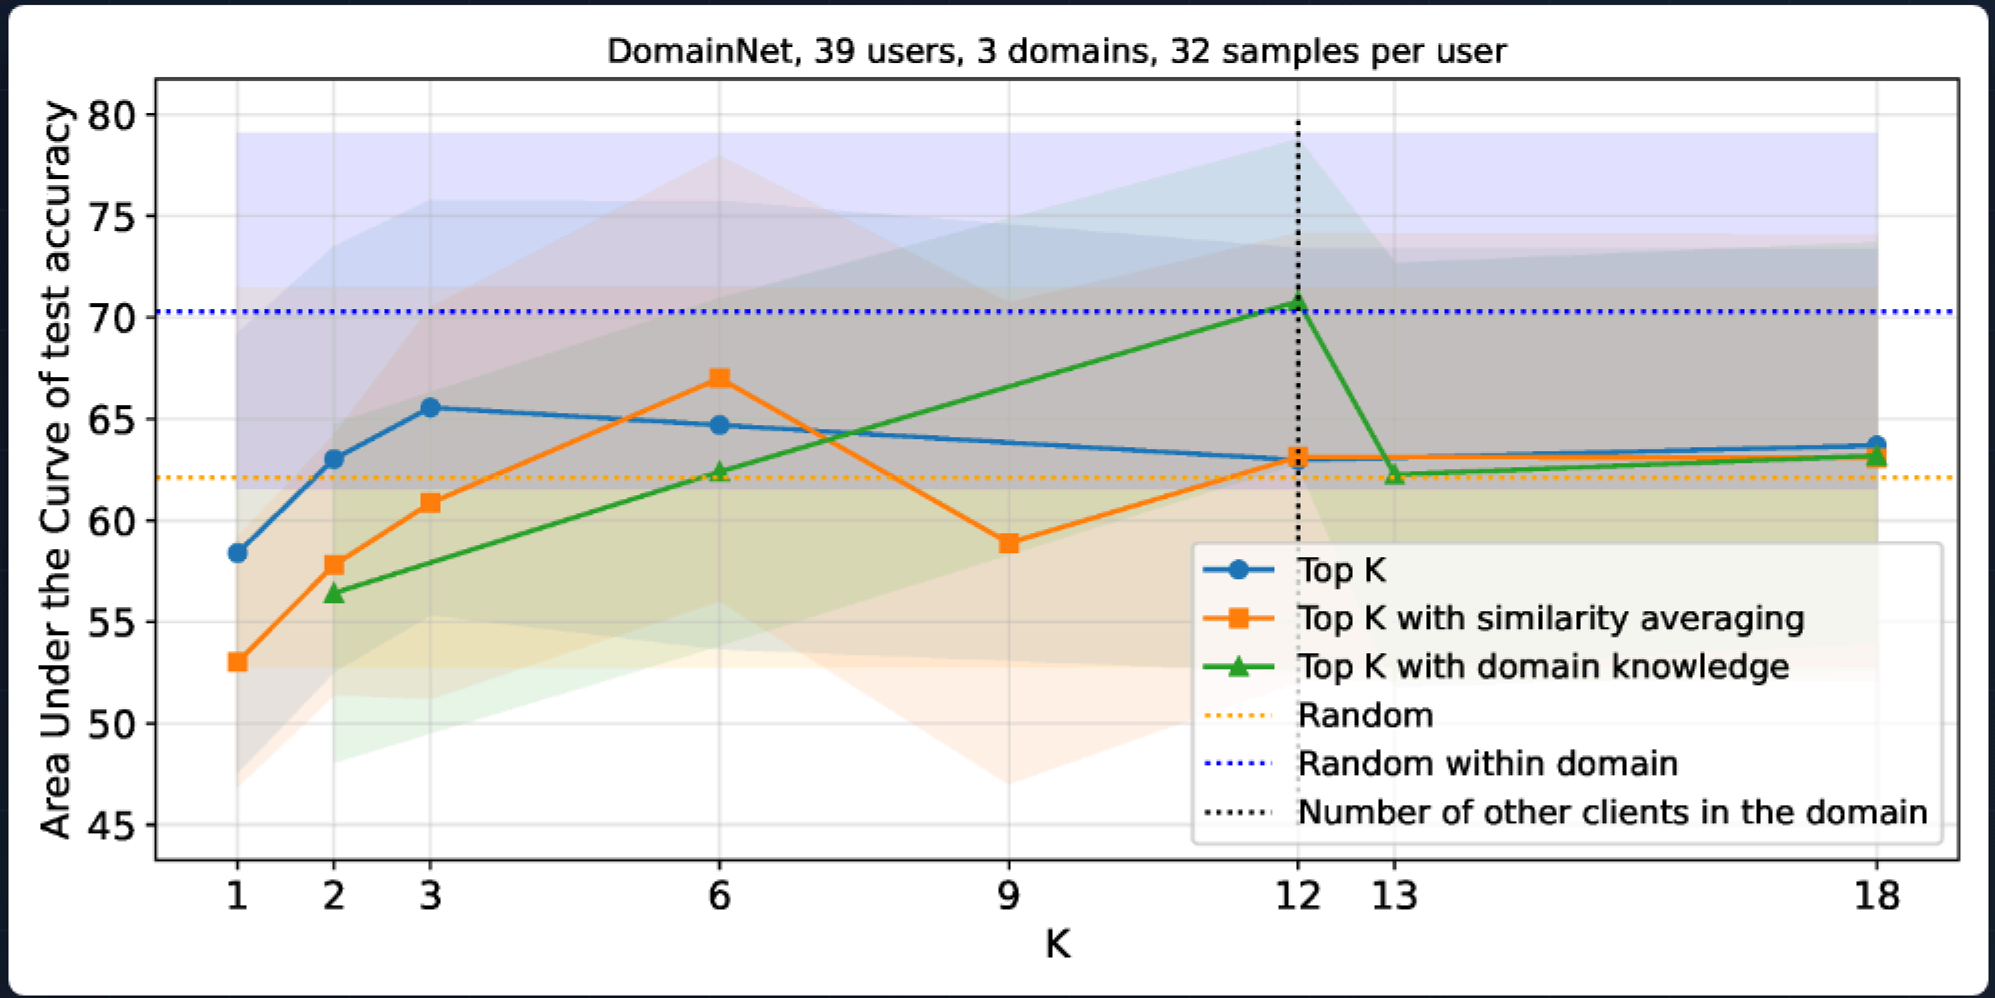

K-sensitivity: optimal K≈12 (domain size) balances diversity and similarity; too small→collapse, too large→cross-domain mixing

Use arrow keys or scroll to navigate

Finding 4: Topology Determines Attack Resilience

→ **Sparse structured graphs more resilient:** Ring/torus degrade gracefully; dense/random show sharp drops

→ **~20% performance difference** between best and worst topologies as attackers grow (0→1→4 malicious nodes)

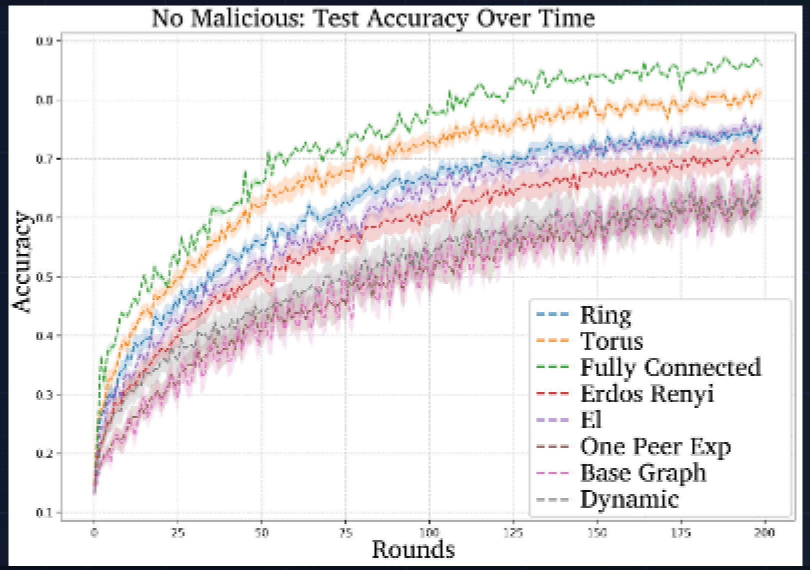

Baseline (0 attackers)

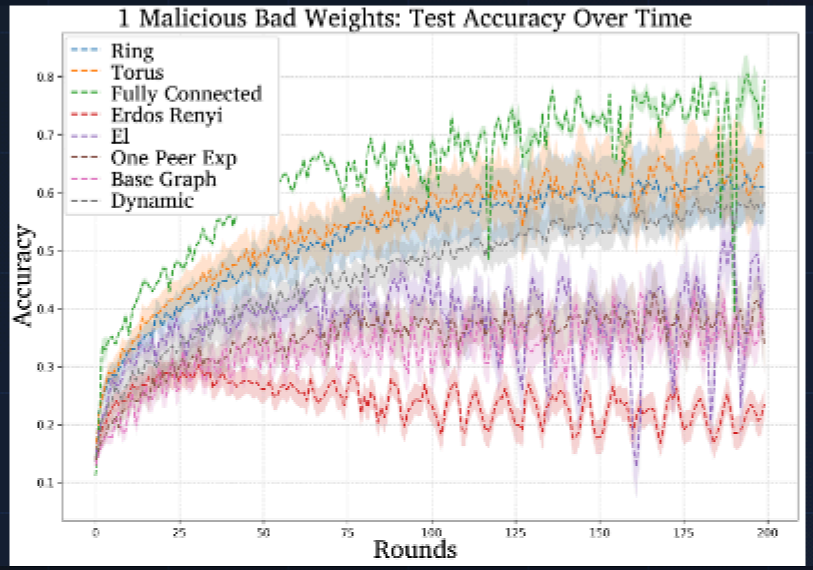

1 malicious node

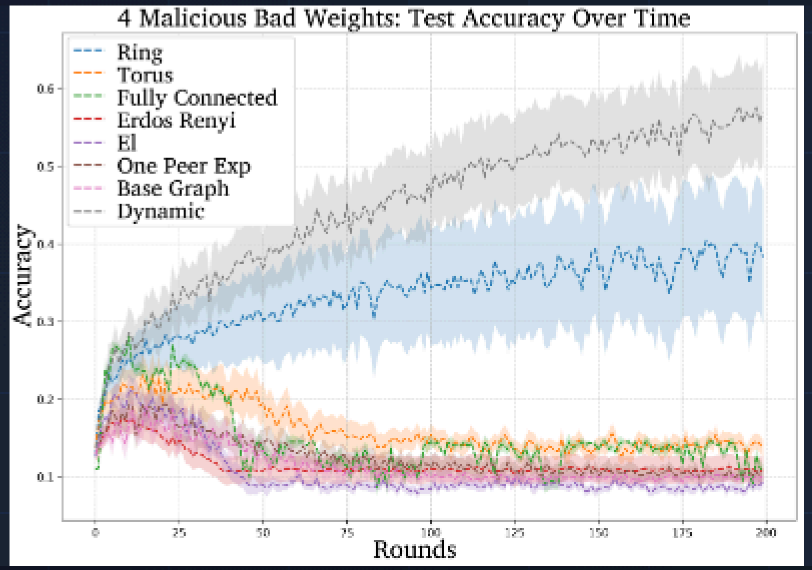

4 malicious nodes

Key Takeaways

- → **Topology is a first-class systems variable** that influences performance, robustness, and privacy
- → **Impact depends on context:** Participant scale, data heterogeneity, communication budget jointly determine effectiveness
- → **Sparse can win:** Structured topologies (ring, torus) achieve comparable/better accuracy than dense graphs at lower cost
- → **Collaborator collapse is real:** Adaptive similarity-based selection can systematically degrade generalization
- → **Decentralized learning = coupled system:** Must jointly reason about optimization AND communication structure

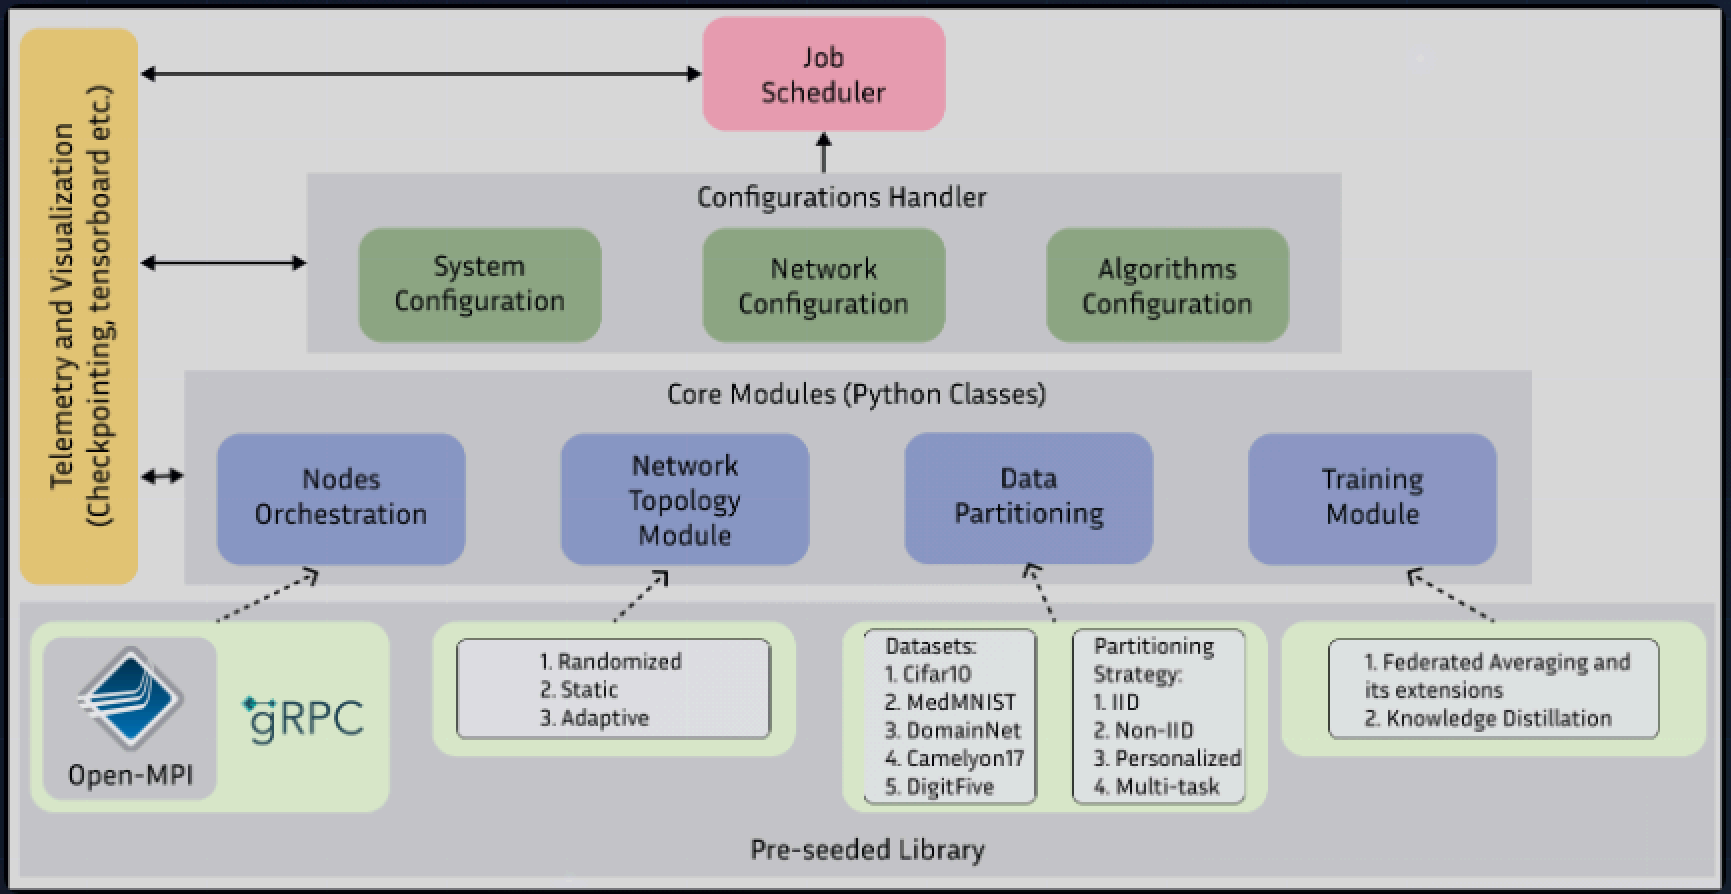

SONAR framework exposes topology as controllable dimension

Conclusion & Future Work

Effective decentralized system design requires jointly reasoning about **topology**, **data distribution**, and **resource constraints**

Impact

SONAR bridges ML and distributed systems: First unified platform for systematic, reproducible topology-aware evaluation

Future Directions

Scaling to larger systems, stronger privacy guarantees, extending to language/multimodal models

Open-source at:

`github.com/aidecentralized/sonar`

Democratizing machine learning through truly decentralized collaboration

Thank you! Questions?